OBJECTIVITY AND BIAS

GORDON BELOT

Abstract. The twin goals of this essay are: (i) to investigate a family of cases in which the goal of guaranteed convergence to the truth is beyond our reach; and (ii) to argue that each of three strands prominent in contemporary epistemological thought has undesirable consequences when confronted with the existence of such problems. Approaches that follow Reichenbach in taking guaranteed convergence to the truth to be the characteristic virtue of good methods face a vicious closure problem. Approaches on which there is a unique rational doxastic response to any given body of evidence can avoid incoherence only by rendering epistemology a curiously limited enterprise. Bayesian approaches rule out humility about one's prospects of success in certain situations in which failure is typical.

Earlier versions of this paper were presented in Abu Dhabi, Ann Arbor, Dubrovnik, Fort Wayne, Los Angeles, New Brunswick NJ, and Rochester—thanks to all those present. For helpful comments and discussion, thanks to several anonymous referees, Frank Arntzenius, Dave Baker, John Bennett, Caro Brighouse, Kevin Coffey, Earl Conee, Daniel Drucker, Katie Elliott, Dmitri Gallow, Jim Joyce, Conor Mayo–Wilson, Alyssa Ney, Laura Ruetsche, Jonathan Schaffer, Chip Sebens, and Jonathan Weisberg. This paper owes a great debt to work of Hilary Putnam, Kevin Kelly, and John Earman.

Forthcoming in *Mind*.

1. Introduction

In the broad, non-judgemental sense in which I will use the term here, a *bias* is a factual or methodological commitment that one brings to empirical inquiry. Biases can stand in the way of our desire to form accurate beliefs in response to evidence. Suppose, for instance, that we aim to determine whether a certain coin is fair from knowledge of the outcomes of a sequence of tosses. If I begin in a state of certainty that the coin is fair and update this opinion in the standard Bayesian fashion, I am guaranteed to maintain my certainty no matter how the tosses turn out.[1] Given my starting point, the only way that I can end up with a true belief is if I start out with one—in effect, my belief that the coin is fair is perfectly insulated from any evidence that I might see.

We want our methods of inquiry to be *objective,* in the sense that they avoid those objectionable biases that tend to undercut our desire to respond to evidence by forming true beliefs. Various notions of objectivity can be defined in terms of the sorts of biases that they permit or prohibit. At one extreme we have:

> Strict Objectivity. A method of inquiry is *strictly objective* if it is entirely free of bias.

This notion is mere fantasy. It is a commonplace that no inductive learning is possible in the absence of substantive expectations about what the world is like.[2] Only against a background of such expectations will I turn to a microscope rather than a telescope or a kaleidoscope if I am interested in the spread of disease. In the limit in which we imagine a scientific *tabula rasa,* we imagine someone who sees every similarity and every dissimilarity as being equally important—and who is therefore unable to form any coherent expectations about the future.

Bias-free inquiry is impossible, then. What sorts of biases are acceptable and which are to be avoided? What sort of objectivity should we aim for? One natural thought is that bias is admissible so long as it washes out in the long run—one is entitled to any starting point, so long as it does not prevent one from arriving at the truth when exposed

[1] Here and below, all claims about responses of Bayesian agents to evidence from coin-tossing experiments follow from the result discussed in Savage 1972 (Sect. 3.6).

[2] On this point see, e.g., Hempel 1966 (Sect. 2.3), Jeffreys 1933 (pp. 524 f.), and Kuhn 1963 (pp. 3 ff.).

to sufficient amounts of the right sort of evidence.[3] Here is one way to make this thought precise.

> Long-Run Objectivity. A method of inquiry is *objective in the long run* for a given problem if, for each hypothesis under consideration, if that hypothesis were true, someone following the method would be more or less guaranteed to eventually settle on the truth (given sufficient amounts of data).

Suppose that there are only three options: the coin is fair, it is so-weighted that it always comes up tails, or it is so-weighted that it always comes up heads. If you are a Bayesian and you initially give each option non-zero credence, then you are more or less guaranteed to end up with credence close to one in the true hypothesis as you see more and more tosses.[4] So your method is long-run objective for this problem—even though it is of course not strictly objective. But if I am still following my method described above, beginning in a state of subjective certainty that the coin is fair and updating by conditionalization, then I remain certain that the coin is fair no matter what evidence I see. So my method fails to be long-run objective: if the coin happens to be so-weighted as to always come up heads, my initial bias prevents me from ever latching on to this fact, even though I will have as my evidence arbitrarily long sequence of outcomes consisting entirely of heads.

There are clearly some problems for which long-run objectivity is out of reach. Suppose, for instance, that we attempt to determine the denomination of a fair coin just from knowledge of whether it came up heads or tails on each of a sequence of tosses. No method can be long-run objective for this problem, because the sort of evidence available fails to distinguish adequately between the relevant hypotheses.

But long-run objectivity can also exceed our grasp even when the evidence available is, intuitively, of the right sort for the problem at hand. Return to the problem of guessing the propensity of a coin to come up heads from knowledge of the outcomes of a sequence of tosses. But suppose now that we countenance each real number between zero and one as a live possibility to be the bias of the coin in favour of heads.

[3] The notion that scientific inquiry is objective in the sense that, properly practised, it is destined to wash out any differences in doxastic and methodological starting points is a perennially popular one. For discussion, see, e.g., Hacking 2000, Hempel 1983, Railton 1994, and Weinberg 2001 (Chs. 12 and 13). For the historical roots of this notion of objectivity, see Daston and Galison 2007 (Ch. V).

[4] That is: no matter which one of these hypotheses is true, in the infinite long run there is zero chance that you will see a sequence of tosses that would frustrate your desire to end up with credence arbitrarily close to one in the true hypothesis.

And suppose further that we are required to announce after each toss which hypothesis we consider most plausible given the data that we have seen so far. Two famous methods for addressing this problem:

> THE STRAIGHT RULE: if you have seen m heads in n tosses, conjecture that the propensity of the coin to come up heads is given by $\frac{m}{n}$.
>
> THE RULE OF SUCCESSION: if you have seen m heads in n tosses, conjecture that the propensity of the coin to come up heads is given by $\frac{m+1}{n+2}$.[5]

Notice that no matter what data these methods see, they always conjecture that the propensity of the coin to come up heads is given by a rational number. But that means that if the true propensity is given by an irrational number, these methods can never arrive at the truth—so *most* hypotheses under consideration have the feature that they could never be discovered by these methods. They fail, in a fairly dramatic fashion, to be long-run objective for this problem.

The difficulty is not special to the straight rule and the rule of succession. In the present context, we can think of methods as being (or, at any rate, as determining) functions that take as input a finite data set and give as output the hypothesis that they consider most plausible given that data. The set of finite data sets for our problem is countable (we can enumerate it—H, T, HH, HT, TH, TT, and so on) while the space of hypotheses is uncountable. There can be no function from a countable set to an uncountable set with the feature that every member of the latter is the image under the function of some member of the former: countably many members of the target set will be the image of some member of the given countable set, uncountably many will not be. But this is just to say that for the problem at hand, for every method, most hypotheses have the feature that they will never be conjectured by that method—no matter how much and what kind of data it sees. And should such a hypothesis be the true one, the method in question will never arrive at the truth, no matter what data it sees. So long-run objectivity fails dramatically for every possible method.[6]

But surely the problem of determining the propensity of a coin to come up heads is one that can be handled objectively, if any nontrivial

[5]This rule, due to Laplace, may look strange—but it is the rule followed by Bayesians who start out with flat priors over the space of possible biases of the coin. For its history and justification, see Zabell 1989.

[6]More generally, long-run objectivity fails whenever the cardinality of the space of data sets is smaller than the cardinality of the space of hypotheses. But this sort of mismatch of cardinality is not the only obstruction to long-run objectivity—see fn. 60 below.

problem can be. Even though we cannot have methods that will eventually fix upon the true propensity-for-heads, no matter what it should be, there are plenty of methods with the feature that they will output a sequence of conjectures more or less guaranteed to converge to the truth in the long run—the straight rule and the rule of succession being among the most simple-minded such methods.

> Asymptotic Objectivity. A method of inquiry is *asymptotically objective* for a given problem if, for each hypothesis under consideration, if that hypothesis were true, the beliefs of someone following that method would be more or less guaranteed to converge to the truth (given sufficient amounts of data).

Something in the area of this notion plays a prominent role in our thinking about objective inquiry.

The twin goals of this essay are: (i) to investigate a family of cases in which even asymptotic objectivity is beyond our reach; and (ii) to consider the implications of such cases for three sorts of approaches prominent in contemporary epistemological thought (each of which can be thought of as taking a stance on questions of objectivity and bias). The point will not be to criticize these approaches for failing to underwrite methods that solve intractable problems—that it cannot underwrite a method guaranteed to converge to the truth is no strike against your favourite approach if *no* method comes with such a guarantee. The point rather is that each of the three epistemological strands has undesirable consequences when confronted with the existence of such problems. Approaches that follow Reichenbach in taking asymptotic objectivity to be the characteristic virtue of good methods face a vicious closure problem. Approaches on which there is a unique rational response to any body of evidence can avoid incoherence only by fencing epistemology into a very small plot. Bayesian approaches rule out humility about one's prospects for success in certain cases where failure is typical.

Section 2 sets up the sort of problems that we will focus on. Sections 3–5 develop these problems and explore their implications for our three strands of epistemology. In section 6 we return to the question of which forms of bias are admissible, drawing some inconclusive conclusions. An appendix is devoted to showing that the difficulties we turn up are not artifacts of the special sort of space of hypotheses that is in play in sections 2–5.

2. Curve Fitting!!

Here is a highly idealized picture of one aspect of the scientific method. One begins with a set of hypotheses, $\mathcal{H}$, concerning the nature of some system. As one gathers data concerning this system, some hypotheses in $\mathcal{H}$ are ruled out. At any stage of inquiry, however, many remain in play. If pressed to select the most plausible one, a scientist will rely on background knowledge, judgements of prior probability, theoretical virtues, methodological principles, favourite statistical tests, and so on. This is of course a stylized picture. But it reflects fairly accurately the workings of some parts of science—especially those concerned with discovering the structure of individual systems rather than discovering laws of nature.

Elementary discussions of the scientific method often focus on a special case of this general picture: curve-fitting. Here is a typical set-up. Physical quantity Y is known to depend only on physical quantity X. A scientist aims to determine which function F encodes this dependence (typically, the scientist will restrict consideration to some special class of hypotheses, such as the space of smooth functions relating X to Y). Data come in the form of ordered pairs (x, y) consisting of a value x of X and the corresponding value $y = F(x)$ of Y. After each data point is revealed, the scientist is expected to make a conjecture: to choose the function in $\mathcal{H}$ that is the most plausible candidate to be F, from among those consistent with the data seen so far (and also perhaps to guess whether F has or lacks some further property of interest). As far as the problem of identifying F goes, the scientist should plot the data points on a piece of graph paper, and select the most plausible candidate for F among those functions whose graphs pass through the data points.

We will focus on a special case of this sort of curve-fitting problem.

(i) X will range over the natural numbers, 1, 2, 3 ....
(ii) Y takes as values only 0 and 1.
(iii) The space of hypotheses, $\mathcal{H}$, is the space of all functions from the set of natural numbers to the set $\{0, 1\}$. There is a natural correspondence between functions of this form and infinite binary sequences: $F \leftrightarrow (F(1), F(2), F(3), \ldots)$. It will normally be convenient to think of our hypotheses as sequences rather than functions.
(iv) The 'curve-fitter' is shown the bits of the true hypothesis one by one in the obvious order—first $F(1)$, then $F(2)$, etc. (Think here of F as a rule that deterministically generates the data points.)
(v) There are two sorts of problem a 'curve-fitting' agent may face.

(a) The Identification Problem: to guess which sequence is being revealed bit by bit.
(b) The Property Problem: for a given property, to guess whether or not the sequence being revealed has that property.

After each new bit is revealed, agents facing the Identification Problem must offer a conjecture as to the true identity of the sequence being revealed and agents facing a Property Problem must offer a conjecture as to whether the sequence has the property in question.

A method for addressing a problem of either sort is a function from the set of possible finite data sets (the set of finite binary strings) to the set of possible answers to the relevant question (sequences in the case of the Identification Problem, Yes-No answers in the case of a Property Problem).

> Asymptotic Reliability. A method is *asymptotically reliable for a given problem and hypothesis* if: were that hypothesis true (i.e., were the agent shown the sequence in question bit by bit), the method would generate a sequence of conjectures that converged to the truth about that hypothesis (more on convergence in a second). A method that is asymptotically reliable in this sense for each of the hypotheses in play is called *asymptotically reliable for the given problem.*

This last notion will henceforth be our official explication of the notion of asymptotic objectivity for problems of the above type.

What sort of convergence is relevant for these problems? Informally, to say that a sequence of objects converges to a given object is to say that one can achieve any desired degree of approximation to the given object simply by taking objects that occur sufficiently far along in the given sequence. There are many explications of this informal idea, adapted to different mathematical contexts. For present purposes, we can get by with restricting attention to the case of real numbers: a sequence $(x_1, x_2, \ldots)$ of real numbers converges to the real number x just in case for any $\varepsilon > 0$ there is an N such that $|x - x_k| < \varepsilon$ for all $k > N$.[7] So in order to approximate the target x as well as you like, you

[7] A more natural route to the characterizations arrived at in the next two paragraphs would proceed via the topological notion of sequential convergence: a sequence $(x_1, x_2, \ldots)$ of points in a topological space X converges to a point $x \in X$ just in case for every open set U containing x, there is an N such that $x_k \in U$ for all $k > N$. For the Property Problem, the relevant space would be a two-point space with the discrete topology (according to which every subset of the space is open). For the Identification Problem, it is natural to think of $\mathcal{H}$ as equipped with the

need only look sufficiently far down the sequence $(x_1, x_2, \ldots)$; for any small but finite neighbourhood of x, the x_k eventually all lie in that neighbourhood.

In our Property Problem, the true binary sequence σ is revealed to the agent one bit at a time—and the agent is required to hazard a guess, after each bit is revealed, as to whether σ instantiates some property of interest. Let us encode YES as two and NO as minus two. On the one hand, we have a number encoding the truth about whether σ has the property of interest—two if it does, minus two if it does not. On the other hand, we have a sequence of twos and minus twos encoding the successive guesses regarding this issue hazarded by the curve-fitter after the revelation of each bit of the true sequence. The curve-fitter converges to the truth for the Property Problem just in case the sequence of twos and minus twos encoding these guesses converges to the number encoding the truth about whether σ has or lacks the property of interest. Applying the standard criterion of convergence we find that the sequence of the agent's guesses converges to the truth in the relevant sense if and only if there is a point in time after which the agent permanently gives the correct answer (to see this, consider $\varepsilon < 4$). Example: the following is an asymptotically reliable method for determining whether the sequence being revealed is all zeroes—assume that it is unless and until you see a one.

Consider next the Identification Problem. Let $\sigma \in \mathcal{H}$ be the true sequence being revealed. After each each bit is revealed, the agent puts forward as a conjecture one of the elements of $\mathcal{H}$—let σ_1 be the conjecture put forward after the first bit is revealed, σ_2 be the conjecture put forward after the second bit is revealed, and so on. Each element of $\mathcal{H}$ is a binary sequence and so can be identified with a real number between zero and one in a natural way.[8] And so we can again apply the ordinary notion of convergence for sequences of real numbers. Now we find: the sequence of conjectures $(\sigma_1, \sigma_2, \ldots)$ converges to the truth, σ, if and only if for each $k > 0$ there is an N such that if $m > N$ then σ_m has the same first k bits as σ—the agent's conjectures are eventually

product topology induced by thinking of it as the product of countably many such two-point topological spaces. When equipped with this topology, $\mathcal{H}$ is known as the *Cantor space.* For each finite binary string w, let B_w be the set of all sequences that begin with that string. Each B_w is an open subset of the Cantor space; and each open subset of the Cantor space is a union of such sets.

[8] Usually it is fine to think in terms of the 'obvious' correspondence $\sigma = (x_1, x_2, \ldots) \mapsto \sum_k \frac{x_k}{2^k}$. But this map has its pathologies—e.g., it counts both $(1, 0, 0, \ldots)$ and $(0, 1, 1, \ldots)$ as corresponding to $\frac{1}{2}$. The map $\sigma \mapsto \sum_k \frac{2x_k}{3^k}$ is better—it is a homeomorphism from the Cantor space to the standard ternary Cantor set.

right about arbitrarily long initial segments of σ.[9] (Compare: if a sequence $(x_1, x_2, \ldots)$ of real numbers converges to π, then from a certain point onwards in this sequence, if we look at decimal expansions, each x_k begins with a 3, and from a certain point onwards each x_k begins with a 3.1, and from a certain point onwards, each x_k begins with a 3.14, and so on—and if this fails, then the sequence does not converge to its target.)

That, then, is going to be our model of inquiry: agents facing the Identification Problem or versions of the Property Problem as a binary sequence is revealed to them one bit at a time. This model constitutes the simplest possible example of a setting in which the supply of data in principle available is never exhausted at any finite stage of inquiry. It is this feature—along with a certain more technical condition—that drives the considerations developed below.[10] So the problems that we encounter will also bedevil more interesting models of empirical inquiry (see the appendix for an example that is more obviously related to real science).

Before proceeding, it is worth noting a couple of features of the Identification Problem our curve-fitters face. (i) As is common in discussions of curve-fitting, we are in effect assuming that our data are error-free. This is unrealistic. But it should strengthen rather than weaken the prospects of the epistemological approaches under consideration below.[11] (ii) In our model, every false hypothesis is eventually ruled out definitively by the data (since any false hypothesis would have to disagree at some point with the true sequence—and eventually this disagreement will be revealed to our agent). This too serves to stack the deck in favour of the approaches under consideration.

3. Reichenbachian Approaches

Let us for the time being take as our fundamental unit of evaluation: a method of inquiry in application to the problem of determining which

[9] To see this, note that each of the sets B_w defined in fn. 7 above is open and apply the topological criterion of sequential convergence.

[10] The technical condition: that certain ways of constructing data streams correspond to hypotheses in the relevant space of hypotheses; see Step (v) in the proofs of Facts 1 and 6 below (and also fn. 61).

[11] There are approaches to making sense of inductive inquiry that turn crucially on the fact that real data are not perfectly accurate: see, e.g., Kemeny 1953 or Forster and Sober 1994. I do not here have space to address these interesting approaches. In brief: I believe that they ultimately push the bump under the carpet to another, very interesting, corner of the same room.

of a set of mutually exclusive (but not necessarily exhaustive) alternatives obtains. Sipping and tasting is a good method for distinguishing between water and wine, a lousy method for distinguishing between water and heavy water, an absolutely hopeless method for determining whether one is a brain in a vat.[12]

It is natural to view a tendency to lead to the truth as more and more evidence is seen as a desirable feature of methods of inquiry. And it is also natural to look askance at methods that lack this feature—if we know that our method may well not be tracking the truth, then we know that the beliefs we form, even after seeing large bodies of evidence, may well be determined by our starting point rather than by the world. These reflections ought to get one in the mood to take asymptotic objectivity as the key epistemic notion, if anything will.

Consider again the straight rule and the rule of succession (discussed in Sect. 1 above) applied to the problem of determining the bias of a coin from knowledge of the outcomes of a sequence of tosses.[13] These methods are (essentially) guaranteed to output sequences of conjectures that converge to the truth as more and more data is seen. Of course, for some sequences of tosses, the straight rule will perform better than the rule of succession—and for other sequences of tosses, vice versa. If all we care about is convergence to the truth, then there is no sense in which one rule is better than the other in overall performance.[14]

In its pure form, *Reichenbachianism* consists of two theses.[15]

(R1) A method of inquiry counts as a good approach to addressing a given problem if and only it is asymptotically objective for that problem.

(R2) One is justified in taking a first-rate doxastic or pragmatic attitude towards the outputs of a method of inquiry applied to a given problem just in case it is good in this sense.

[12]Note, though, that habitual guzzling can be part of a good method for distinguishing between water and heavy water—see Leslie 2013 (p. 151).

[13]Structurally similar points could be made using various methods for handling traditional curve-fitting problems, such as piece-wise linear interpolation (connect-the-dots) and Lagrange interpolation (use the lowest-degree polynomial consistent with the data). See the discussion of Hempel 1966 (Sect. 4.4).

[14]For this sort of point, see Reichenbach 1949 (Sect. 91).

[15]For canonical presentations, see Reichenbach 1933, 1938 (Ch. V), and 1949 (Sect. 91). Elements of Reichenbach's account can be found in, e.g., Anderson 2004, Earman 1993, Enoch and Schechter 2008, Feigl 1934 and 1954, Kelly and Glymour 1989, Kelly 1996, and Wright 2004. The label is of course somewhat arbitrary. Reichenbach doesn't always sound thoroughly Reichenbachian. And ideas similar to Reichenbach's can be found already in Peirce: see Reichenbach 1939 (pp. 187 ff.) and Madden 1964 (pp. 132 ff.).

Pure Reichenbachian approaches differ from one another concerning which attitudes towards theories are involved; there are also exist less pure strains of Reichenbachianism that take asymptotic objectivity to be necessary but not sufficient for goodness.[16]

Every relatively pure Reichenbachian faces some awkward facts.[17] If a problem admits an asymptotically objective method, then it admits many such methods: the straight rule and the rule of succession are just two of many asymptotically objective methods for determining the bias of a coin; likewise, there are many asymptotically objective techniques for standard curve-fitting problems. Worse, because asymptotic objectivity is a condition on the behaviour of a method in the infinite long run, methods that respond in bizarre ways to data sets below some fixed size can still be asymptotically objective. In short, it is natural to worry that Reichenbachianism is far too permissive in drawing the line between acceptable and unacceptable methodological biases.

Perhaps one can learn to live with this—if, for instance, one thinks of Reichenbach's approach as embodying the strongest interesting thing one can say in response to Hume's problem of induction. My goal here is to develop another sort of objection—one that is, in my own experience, more difficult to live with.

Consider a pure Reichenbachian account in the context of the problem described in the preceding section. Here the account tells us to (say) believe our conjectures if and only if our method is asymptotically reliable for the problem at hand.

Consider first the Identification Problem, in which our agent is asked to guess the identity of the sequence being revealed. Asymptotically reliable methods for this problem are easy to come by: so long as at each stage one conjectures a sequence that is consistent with the data seen so far, one's sequence of conjectures will converge to the true sequence (since from the kth stage onwards, each conjecture will be correct about the first k bits of the true sequence). Here are two simple-minded methods that do the trick for this problem:

> Method A: conjecture the sequence that results from tacking an infinite sequence of alternating zeroes and

[16] It is natural to read Reichenbach 1938 (Sect. 43) as counting a method as good just in case it is asymptotically objective and as taking us to be justified in acting on the outputs of good methods but not others. It is natural to read Earman 1993 (Sect. 11) as requiring further that good methods be Bayesian—and as advising Bayesians to take a non-alienated attitude towards their credences just in case the methods determined by their priors are asymptotically objective.

[17] For further discussion and references, see Salmon 1991. For further objections special to probabilistic contexts, see Hacking 1968 and Sober 1988.

ones on the end of the data seen so far.

> METHOD B: conjecture the sequence that results from tacking an infinite sequence of zeroes on to the data seen so far.

Our Reichenbachian approach will consider both of these to be good methods in virtue of their asymptotic reliability—and so will advise anyone following these methods to believe their conjectures.[18]

Consider next what happens if we further confront our agent with a version of the Property Problem involving a property of the following special sort.

> SLIPPERY PROPERTIES. A property P of infinite binary sequences is called *slippery* if, for any finite data set, among the sequences that extend this data set, some have P while others lack it.[19]

Lots of interesting properties of sequences are slippery in this sense: being eventually constant (all zeroes or all ones from some point onwards); being periodic; encoding a binary expansion of a rational number; being a sequence in which the limiting relative frequency of 0's exists and is equal to $\frac{1}{2}$. The main arguments below go through for any slippery property. For purposes of illustration I will often use the last-mentioned one, which I will call the property of being *fifty-fifty.*[20]

What happens if we ask our agents to guess whether the sequence being revealed is fifty-fifty as well as guessing its identity? If you are following Method A, you had better always guess that the true sequence is fifty-fifty—after all, every sequence that you will ever conjecture in response to the Identification Problem has this feature, and

[18]It may be surprising that the identification problem can be handled so easily. In this regard note: (i) that in the case of more traditional curve-fitting problems, too, asymptotic objectivity is often fairly easily achieved (see, e.g., Hempel 1966, Sect. 4.4); (ii) working in the Bayesian framework, it is easy to construct priors on $\mathcal{H}$ that, when conditionalized on data sets for our problem, lead to posteriors that become arbitrarily tightly peaked about the true hypothesis (Belot 2013a, fnn. 31 and 32).

[19]Equipping $\mathcal{H}$ with the topology described in fn. 7, this is equivalent to: the set of sequences with P and the set of sequences without P are both dense in $\mathcal{H}$.

[20]Each of the properties mentioned in this paragraph corresponds to a subset of $\mathcal{H}$ that happens to be small. Indeed, each is meagre as a subset of $\mathcal{H}$ (for this notion, see fn. 53 below): each countable set is meagre; for an argument showing that the set of fifty-fifty sequences is meagre, see Oxtoby 1980 (p. 99). But that is inessential. We could just as easily work with a subset of $\mathcal{H}$ such that neither it nor its complement is small—e.g., the set of sequences that either begin with a zero and are fifty-fifty or begin with a one and are not fifty-fifty.

because Method A is asymptotically reliable you believe those conjectures. Likewise, anyone following Method B ought to always guess that the true sequence is *not* fifty-fifty.

Although though Methods A and B are asymptotically reliable for the Identification Problem, they lead to unreliable methods for our slippery version of the Property Problem. For suppose that the true sequence is not in fact fifty-fifty—suppose, for instance, that the true sequence is $(1, 1, 1 \ldots)$. At each stage of inquiry, the follower of Method A will guess that the true sequence is fifty-fifty—and the sequence Yes, Yes, Yes, ... does not eventually settle down to No. Similarly, if the true sequence is something like $(1, 0, 1, 0 \ldots)$ then the follower of Method B will put forward the sequence of conjectures No, No, No, ... And this fails to settle down eventually to the true answer, Yes.

Our pure Reichenbachian approach will tell followers of Method A or Method B that they should distance themselves from the output of their method for guessing whether or not the sequence they are being shown is fifty-fifty. And it is easy to think that this is the right thing to do. After all, the followers of Method A in effect begin inquiry in effect certain that the sequence is fifty-fifty—and no evidence that they could see would change their minds. And the followers of Method B are biased just as strongly in the opposite direction. Methods that thus insulate a belief about a matter under investigation from any possible evidence are not ways of finding out about how the world is.

Note where this leads: our Reichenbachian approach is advising people to work with a notion of belief that is not closed under logical implication. The followers of Method A are advised, at each stage of enquiry, to adopt a Grade-A doxastic attitude towards their guess as to the sequence being revealed but to distance themselves from their conjecture that this sequence is fifty-fifty—but at each stage, the sequence that they put forward as being the most plausible candidate to be the true one will in fact be one that is fifty-fifty. It would be deeply bizarre—and not merely funny or a bit odd, as would be the case if it involved a mere violation of pragmatic constraints—if a scientific article were to argue that a detailed model was well-supported by the evidence, but then went on to warn the reader not to leap to the conclusion that the evidence offers the same sort of support for the claim that the true model belongs to a family containing the given model. Imagine that in announcing his discovery that gravitational attraction varies as the inverse square of distance, Newton had cautioned against leaping to the conclusion that it varies inversely as some polynomial or other of distance!

Of course, some epistemologists hold that knowledge or justified belief fail to be closed under logical implication. Advocates of this view advertise it as the lesser of two evils—with the greater evil being skepticism about knowledge or justification (see, e.g., Dretske 2005a,b). Perhaps in our case, too, denying closure is the least evil path. But it is important to note that skepticism is nowhere in view in the present context. We ought to be happy to say that at the fine-grained level none of the hypotheses in play counts as a skeptical scenario: we are interested in the long-run behaviour of methods in a setting in which each false hypothesis is eventually conclusively falsified. The proposition that the true sequence is fifty-fifty is just a disjunction of such fine-grained hypotheses—and we do not normally think that a disjunction of non-skeptical scenarios should count as a skeptical scenario.[21]

Nor should the sort of failure of closure that we see here be assimilated to that which occurs on metaontological accounts in the tradition of Carnap 1950. Perhaps it is a good idea to say that questions internal to the thing framework (such as whether the animal you hear is Grisbi the dog) admit of empirical confirmation and disconfirmation even while maintaining that questions external to this framework (such as whether things exist) have a different status. But in the case at hand, the analog of an external question might be something like the question whether our chosen space of hypotheses is a suitable one—the question whether the true sequence is fifty-fifty is no more external to our framework than the question whether the animal you hear is a dog is external to the thing framework.

So I do not think that this closure problem can simply be brushed aside. But it is still natural to suspect that the problem is an artifact of our focus on Methods A and B. These are ridiculous methods—for instance, no matter what evidence they see, neither will ever conjecture that the true sequence is $(1, 1, 1, \ldots)$. Anyone facing our Identification Problem would use a method incomparably more subtle than these methods—and it is natural to hope that methods one might actually use would be immune to the closure problem developed here, so that the moral of all of this would just be that we should be looking for impure forms of Reichenbachianism that don't count Methods A and B as good. This is a beguiling hope but it is a false one.

[21]Note that Reichenbachian approaches are closely related to a standard response to skepticism, contrastivism about knowledge—and contrastivists take it for granted that knowledge is closed under weakening (see, e.g., Schaffer 2006, p. 262).

Fact 1. There is no asymptotically reliable method for determining whether the true sequence is fifty-fifty. (And likewise for other slippery properties.)

Proof. Suppose for *reductio* that M is an asymptotically reliable method for this problem.[22]

(i) M will have to settle down after a finite time to the correct answer, whichever hypothesis is generating the data seen: the only sequences of YES-NO answers that converge to YES are those that are YES permanently from some point onwards; and likewise for NO.

(ii) For this problem, any finite data set is consistent with either answer to the question being asked.

(iii) It follows from the preceding that for any finite data set, no matter which answer our asymptotically reliable M gives when shown that data, there are ways of continuing the data set that would make M change its answer.

(iv) So we can construct a data stream of bits that will make M flip-flop *ad infinitum* on the question being asked.

(v) And this data stream corresponds to a binary sequence in our space of hypotheses (since it includes *all* binary sequences).

(vi) But of course when shown the data generated by this hypothesis, M never settles on the true answer to the question asked.

This contradicts our initial assumption that M was asymptotically reliable for this problem. So no such M can exist. $\square$

So *every* method for handling this sort of problem faces the same sort of closure problem that bedevils Methods A and B. Such closure problems cannot be escaped by moving to more restrictive impure forms of Reichenbachianism—unless they are so restrictive that they do not recognize *any* good methods for identifying which binary sequence our agent is being shown (which would not leave very much of the basic approach intact). And, in any case, we will see in the appendix below that the same sort of closure problem arises in other settings as well.

The knock against Reichenbachianism has always been that it gets us too little—sure, it counts our method as good, but it also debases it by counting as good many other methods, including some pretty crazy-looking ones. The response to this has always been that one cannot always get what one wants, but one can sometimes get what one needs. The closure problem under discussion shows that Reichenbachianism delivers even less than people have thought. In scientific contexts our

[22] The following is essentially the proof of Prop. 12 in Kelly and Glymour 1989.

method (whatever that is) presupposes that warrant is closed under logical consequence. So the closure problem unearthed above shows us that Reichenbachians are not able to vindicate our method after all.

4. Uniquarianism

Epistemologists are divided on the status of the following thesis.

> Uniqueness: there is a unique rational doxastic response to any given body of evidence (where 'doxastic' should be understood broadly, to cover both full belief and degrees of belief).

Permissivists deny it, *uniquarians* uphold it.[23] Here is the best motivation for uniquarianism.[24] Epistemology should be adequate to the jury room as well as to debates about the objectivity of science. Presumably few of us would vote to convict unless we believed the defendant to be guilty. But if there are multiple rationally acceptable responses to a given body of evidence, then it seems that no matter how responsible one is in weighing the evidence, factors irrelevant to the question at hand (upbringing, neurochemistry, etc.) must be playing a role in determining one's beliefs. And it can be hard to see how one's beliefs concerning serious questions can survive reflection on the fact that they depend in this way on what are essentially chance matters.[25] This would seem to give us some reason to reject permissive approaches in favour of uniquarianism.[26]

The classic worry about uniquarianism derives from the fact that uniquarians have made so little progress towards identifying the principles of universal rationality.[27] Here it would be hard here to top Ramsey's blistering criticism of Keynes. A taste:

[23]Apologies for the neologism. I know of no standard term for partisans of Uniqueness. There exists already an English word *uniquity,* standing to *unique* as *antiquity* stands to *antique.* Here the analogy is extended: *uniquarian* stands to *uniquity* as *antiquarian* stands to *antiquity.*

[24]Uniquarian views can be found in, e.g., Carnap 1945, Feldman 2007, Keynes 1921, Maher 2006, White 2005 and 2009, and Williamson 2000 (Ch. 10). In later writings, Carnap adopted an agnostic attitude on this question: see, e.g., Carnap 1962 (p. 316); for further discussion and references, see Skyrms 1991 (Sect. 6).

[25]Reflections of this sort unsettle thoughtful Reichenbachians: the threat of the alienation from one's own epistemic norms is a common theme among Reichenbachian epistemologists. See, e.g., Reichenbach 1938 (Sect. 43) and 1949 (Sect. 91) and Wright 2004.

[26]The case is put forcefully in White 2005; see also Feldman 2007. For quite different reactions to the same sort of considerations, see Rosen 2001 and Schoenfeld 2014.

[27]Another worry, concerning the rationality of scientific disagreement, will be discussed briefly in Sect. 5 below.

> But let us now return to a more fundamental criticism of Mr Keynes' views, which is the obvious one that there really do not seem to be any such things as the probability relations he describes. He supposes that, at any rate in certain cases, they can be perceived; but speaking for myself I feel confident that this is not true. I do not perceive them, and if I am to be persuaded that they exist it must be by argument; moreover I shrewdly suspect that others do not perceive them either, because they are able to come to so very little agreement as to which of them relates any two given propositions. (Ramsey, 1931, p. 161)

My aim here is to add another sort of worry.[28]

Suppose, for the time being, that uniquarians accept that two agents facing the same data points have the same evidence relevant to the Identification Problem and our Property Problem. The uniquarian thesis that there is only one rational doxastic response to any given body of evidence then implies that rationality determines a function from possible data sets to conjectures about the true identity of the sequence being revealed and as to whether or not it is fifty-fifty. In the terminology used above, rationality determines a method for addressing the Identification Problem and our Property Problem. I make the following claims about this picture.

(1) *The unique rational method is in some sense optimal.* Otherwise, why care about rationality?

(2) *The range of asymptotic reliability of a method for addressing a problem counts towards its optimality.* That is, if all we are told about two methods is that one is asymptotically reliable for a strictly larger subset of our space of hypotheses than is the other, then that gives us reason to prefer the former to the latter. Otherwise, it is difficult to see how to do justice to our sense that part of the point of being rational is that it is, in general, the best way to end up with true beliefs.[29]

[28]In effect, what follows is a generalization and simplification of Putnam's argument against (his conception of) Carnap's program in inductive logic. The role played by recursion-theoretic results in Putnam 1963a,b is played here by elementary facts about slippery properties.

[29]It might seem that someone like Williamson, for whom there is a distinguished prior that 'measures something like the intrinsic plausibility of hypotheses prior to investigation' (2000, p. 211), can afford to deny this—since it seems that all constraints on rationality are loaded into the input end of things on this sort of view. But I doubt that Williamson himself would be willing to take this sort of

(3) *In measuring optimality, asymptotic reliability is not to be traded off against other virtues.* Consider, by way of illustration, a uniquarian position on which in determining the one true method, we should trade reliability off against tractability of implementation.[30] To keep things simple, let us imagine that these are the only factors that are relevant—so we are imagining a uniquarian who maintains that the one true method is the method that achieves an optimal balance of reliability and tractability. This is an untenable position. Uniquarians hold that there is one rational response to a given body of evidence. Their notion of rationality is species- and world-independent. But a notion like 'ideal balance of reliability and tractability' is going to be species-dependent: my dog is not willing (or able) to trade very much tractability for even quite a bit of reliability; I am able to do more—but presumably as I stand to my dog, so some (actual or possible) species of brainiacs stand to me, some species of super-brainiacs stands to them, and so on without end. Each species has an exchange rate at which its members are willing to trade a gain in reliability for a loss in tractability. It is incredible that the species-independent notion of rationality should be arrived at by plugging the exchange rate corresponding to some particular actual or possible species into the formula:

rational = ideal balance of reliability and tractability.

Indeed, suppose, for definiteness, that the rationality-defining rate corresponds to that of brainiacs. Then, of course, humans will, very, very, often be irrational because they are not following a method that they are incapable of following, while super-brainiacs will very, very, often be irrational because they do not live down to the standards of brainiacs. Why should such beings of either sort care about rationality?

A related point. Those uniquarians who are objective Bayesians typically assume that rationality requires logical omniscience: the degrees of belief of rational agents satisfy axioms of the probability calculus, which means that such agents always assign probability one to logical truths. Since the decision problem for first-order logic is undecidable, there is a sense in which Bayesian agents have

step—see his remarks (2000, Sect. 8.7) on rationality as 'a sub-goal on the way to truth'. Of course, some philosophers do deny that there is any close connection between rationality and truth; for discussion and references, see Hookway 2007.

[30] Couldn't there be virtues immune to the sort of argument given here? It is difficult to imagine what they would be like—certainly virtues like short-term reliability, naturalness, and rate of convergence all seem to be relevantly similar to tractability.

computational resources that outstrip the classical models of computation.[31] Bizarre, then, to take tractability into account at all—it is as if, having been handed a magical computer with infinite capacities, we then cavil at using up disk space.

Uniquarians who accept the foregoing are in trouble. For they are committed to the view that rationality determines an optimal method for approaching our Property Problem, where optimality implies maximal asymptotic reliability (i.e., there can be no method asymptotically reliable on a set of hypotheses larger than the set of hypotheses on which the rational method is asymptotically reliable). But there is no maximally asymptotically reliable method for the problem of determining whether or not the sequence being revealed is fifty-fifty.

Fact 2. For any method of forming conjectures as to whether the true sequence is fifty-fifty, there is another method that is asymptotically reliable on a strictly larger subset of the space of hypotheses. (And likewise for other slippery properties.)

Proof. Let M be any method. And let σ_0 be a hypothesis with respect to which M is not asymptotically reliable. Suppose for definiteness that σ_0 is not fifty-fifty. Define a new method M^* as follows: if fed a data set consistent with σ_0, M^* conjectures that the true sequence is not fifty-fifty; otherwise M^* mimics the conjecture made by M on the same data. If the true sequence is not σ_0, then from a certain point onwards M^* will always make the same conjecture as M. So M^* is at least as asymptotically reliable as M. But if the true sequence is σ_0, then M^* always makes the correct conjecture—so unlike M, M^* is asymptotically reliable for σ_0. So M^* is asymptotically reliable for a strictly larger set of hypotheses than is M. $\square$

Uniquarianism is incoherent—if it involves (1)–(3) above and condones the identification of an agent's evidence with the data seen.

But wait! This accusation of incoherence is easily avoided: the denial of the supervenience of evidence on data is built into some prominent uniquarian accounts. Some hold that evidence cannot supervene on data because when we consider two different situations in which a given agent might face a given body of data, we in general find differences in the evidence that the agent would possess. And some hold that evidence cannot supervene on data because when we consider two

[31]Well, that is too quick—in principle, computable Bayesians can assign probability one to each logical truth, so long as they are willing to assign probability one to some contingent formulas. But this approach has very high costs—see Field 2009 (pp. 257 f.).

different agents facing a given body of data, we in general find a difference in the evidence available. Let us consider these two approaches in turn.

Williamson (2000, Ch. 9) defends an account of the first kind.[32] On his view, an agent's evidence is the set of propositions known by the agent. Imagine that a given curve-fitter faces a data set consisting of a zillion zeroes. If the true sequence is all zeroes then, presumably, at some point the curve-fitter comes to know that the true sequence is the zero sequence (to deny this would be to deny that induction can produce knowledge).[33] On Williamson's view: in this case, the curve-fitter's evidence will include the fact that all future bits revealed will be zero.[34] But clearly, there will be other cases in which the same curve-fitter facing the same zillion data points lacks that evidence—e.g., because the next bit of the true sequence is in fact a one.

Now, it may well be that the central contention here is correct: what it is rational for an agent to believe does not, in general, supervene on the publicly available scientific evidence because what it is rational for an agent to believe is determined by that agent's evidence—and this in general depends not only on the publicly available scientific evidence but also on further facts external to the agent. But one might well have hoped that in certain cases at least—such as highly idealized toy models of scientific inference—such complications could be set aside so that one could give an account of scientific inference that involved only scientific evidence and that was operational in Williamson's sense—'a set of rules such that one is, at least in principle, always in a position to know whether one is complying with them'.[35] If it is insisted that no such cases exist, then one rescues uniquarianism from the threat of incoherence at the price of severing the connection between epistemology and the many projects in statistics, artificial intelligence, machine

[32]It is also an account of the second kind—indeed, Williamson (2000, p. 210) suggests that it may be impossible for perfectly rational beings to have the same evidence as us in *any* situation.

[33]In conversation, Williamson suggests that whether there is knowledge here may depend on why the true sequence is all zeroes—e.g., whether it is so by law or by fluke. Let us assume that it is so by law in our case.

[34]Actually, there is considerable tension between this consequence of Williamson's view and one of the main arguments he gives for his thesis that Evidence = Knowledge (2000, pp. 200 f.). Thanks to Daniel Drucker for discussion of this point.

[35]Williamson 2008 (p. 277). See also Williamson 2000 (Sect. 8.7). Williamson's arguments against the possibility of operationalizing epistemology turn on Sorites considerations that are inapplicable in the present context.

learning, and confirmation theory that presuppose that there are substantive things to be said about what curve-fitters and their ilk should believe, given the data they have seen.[36]

As I understand them, Conee and Feldman defend an account of the second kind.[37] They distinguish (2008, Sects. 1.1 and 4.2) between *scientific evidence* (that is publicly available and reliably indicates truth) and *justificatory evidence* (that may be private, and is the sort of evidence upon which justification supervenes). Agents may face the same scientific evidence but differ in the justificatory evidence that they possess (2008, Sects. 1.1 and 3.5): two detectives with the same background knowledge and looking at the same clues at a crime scene can differ in their justificatory evidence (and hence also in what they are justified in believing) because only one of them sees that the guilt of a certain suspect would best explain the presence of these clues. Suppose that you and I face the same set of data for the Identification Problem. You are an expert who can just 'see' very early on that the best explanation for the data we are seeing is that the true sequence corresponds to the binary expansion of π while I am a neophyte who cannot—so early on you will have more justificatory evidence than I do, even though we are looking at the same data.[38]

But if the legitimacy of idealizations in which evidence is determined by data is denied on grounds of this sort, we land in a new kind of trouble. For now uniquarianism becomes something like: relative to each data set and each cognitive constitution, there is a unique rational doxastic response. But what does it mean to take my cognitive limitations into account in considering what it would be rational to believe were my evidence to include all quadrillion data points surveyed in the experiment that led to the discovery of the Higgs boson? It seems that the only sensible answer is that since my cognitive limitations render it impossible for me to comprehend such a data set, the unique rationally prescribed response must be some kind of doxastic neutrality, or some other state that will not depend in any interesting way on subtle features of the data. But it does not seem absurd to ask what we

[36]The flip-side is that on an account like Williamson's, activities like curve-fitting becomes at least quasi-mystical endeavours—in that they involve at least the crucial first two of the four marks of mystical experience (ineffability, noetic quality, transiency, and passivity) identified by James 1902 (Lectures XVI and XVII).

[37]It may also be an account of the first kind—Feldman and Conee (2004, Afterword Sect. 2) maintain that it is consistent with their view that Evidence = Knowledge.

[38]Compare with the discussion of Examples 3 and 4 in Feldman and Conee 2004 (Sect. II).

should believe in light of this body of data—and to expect the answer to depend nontrivially on its detailed structure.

Uniquarians must deny, on pain of incoherence, that our agents' evidence can be taken to supervene on the data they have seen, even in the most idealized of settings. Such denial in fact has fairly deep roots in some contemporary forms of uniquarianism. But the price is steep: scientists or others who come knocking, hoping that epistemology might have helpful advice for those puzzling over what we should believe given this or that vast body of data in the contemporary sciences must be sent packing—and told to come back only when they are in a position to also tell us something about the answers to certain substantive open questions that go beyond the data or when they are ready to ask questions about data sets surveyable by human beings. Many will find this price too high—and will prefer an account of epistemology on which there is at least sometimes room to ask what it is rational to believe on the basis of a large body of scientific evidence.

5. Bayesianism

On *Bayesian* accounts of rationality, rational agents can be thought of as beginning life equipped with a probability measure (their *prior*) over some relevant space of hypotheses, and as having their doxastic states at later times determined by conditionalizing this measure on the evidence available to them at that time.[39] Orthodox *subjective* Bayesianism recognizes no substantive constraints on rational priors. Various forms of *tempered* Bayesianism recognize some such constraints—*objective* Bayesianism being the limiting case in which a unique rationally permitted prior is posited. For expository convenience, I will focus on the subjectivist approach—but, as will be clear, my argument applies equally well to tempered and objectivist approaches.[40]

By my lights, the best route to subjective Bayesianism is the following (this may be completely different for others). A standard criticism of uniquarianism is that because it makes rational disagreement between people sharing the same evidence impossible, it makes it difficult to account for the behaviour of scientists during scientific revolutions.[41] Reichenbachian approaches do not have that problem—but they are

[39]For guides to the literature on Bayesianism in its ever-proliferating varieties, see, e.g., Earman 1992, Howson and Urbach 2006, and Joyce 2010.

[40]Since objective Bayesianism is a form of uniquarianism, it also faces the challenge developed in the preceding section. Some generalizations of the basic Bayesian approaches will be briefly discussed at the end of this section.

[41]Howson and Urbach 2006 (Sect. 8.f). Something along these lines is implicit in Kuhn 1963 (Ch. XII) and 1970 (pp. 134 f. and 156 ff.).

generally taken to err in the opposite direction by endorsing all sorts of patently crazy methods. Wouldn't it be nice if we had an account with the flexibility to treat (selected) actual disagreement as rational that also required agents to follow rules that could be given some sort of motivation beyond their asymptotic behaviour? Subjective Bayesianism is an especially simple and attractive account of this kind. In a sense, all rational agents obey the same canon (updating by conditionalization). Nonetheless, agents facing the same evidence need not agree—and in general will not if their priors differ. Famously, the account has an interesting, but not fully compelling, rationale: unless one's instantaneous degrees of belief are representable by a probability measure, one is a potential mark of Dutch bookies; and unless one's credences evolve by conditionalization, one is a potential mark of Dutch bookies.[42] In a sense, subjective Bayesianism blames all rational disagreement between evidentiary peers on initial doxastic differences. They are at pains to argue that the amount of room this leaves for rational disagreement is *just right.* On the one hand: the space of doxastic starting points is rich enough—e.g., Bayesians are capable of providing models of actual scientific behaviour.[43] On the other hand: for a range of problems, Bayesian agents are certain that their posterior probability distributions will become more and more tightly peaked around the true hypothesis (i.e., they assign probability zero to the set of data streams that would frustrate such convergence); and for a (somewhat different) range of problems, Bayesian agents are in fact certain to converge to the truth, provided that their priors are suitably spread out over the entire space of hypotheses in play.[44]

For present purposes, it is convenient to note two of the many objections to subjective Bayesianism.[45]

(i) If all priors are on a par so far as rationality is concerned, then patently irrational priors will count as rational.[46]

(ii) Scientists give arguments intended to have objective force. We want to know whether the evidence adduced *really* supports the conclusions in question. It is of no interest to reconstruct the

[42] There are many worries about each of these pillars; see, e.g., Hájek 2008 and Moss (forthcoming). For discussion of several other important rationales for subjective Bayesianism, see Joyce 2010.

[43] For doubts about this claim, see Belot (forthcoming).

[44] For discussion and references, see Earman 1992 (Ch. 6) or Belot 2013a.

[45] For discussion of further objections, see, e.g., Earman 1992, Howson and Urbach 1989 (Ch. 11), and Joyce 2010.

[46] See, e.g., Boghossian 2008 (p. 423) and Salmon 2005 (Sect. 4). For further discussion, see Joyce 2004 (Sect. 5).

arguments given so that they are about the personal opinions of scientists.[47]

These are often thought to provide grounds for tempering subjective Bayesianism—for placing more or less strong constraints on rationally permitted priors. My goal is to show that consideration of our Property Problem pushes in the opposite direction: it is not (just) that subjective Bayesianism is too liberal, it is (also) too strict, in that it deems certain apparently rational stances to be irrational.

It will take a while to work up to the point. Let us first of all consider a bit more the obstructions to asymptotic reliability that stand in the way of any method for determining whether the sequence being revealed has some particular slippery property.

Let us say that a hypothesis *flummoxes* a method for resolving a YES-NO question if it makes the method flip-flop *ad infinitum* between YES and NO. And let us say that it *fools* the method if it leads the method to eventually settle permanently on the wrong answer to the question asked. Let the *failure set* of a method relative to a given YES-NO question be the set of all hypotheses that flummox or fool it. Let us call a method for determining whether the true sequence has some given property *open-minded* if it has the feature that no matter what data it sees, there are ways of extending the data set that would make it change its mind regarding this question.[48]

From Facts 1 and 2 above it follows immediately that every method for determining whether the true sequence has a given slippery property has an infinite failure set: for every method, there is another whose failure set is strictly smaller; so if any method had a finite failure set, then there would be a method with an empty failure set; but every method has a non-empty failure set.

By exploiting the fact that methods that are difficult to flummox are easy to fool and the fact that open-minded methods are easily flummoxed, we can show that this is only the first of several senses in which the failure sets of interest are large.[49]

[47]See, e.g., Chalmers 1976 (p. 188), Glymour 1980 (pp. 74 f.), Horwich 1982 (pp. 28 ff.), and Salmon 2005 (Sect. 4).

[48]It is not difficult to find an open-minded prior for the problem of determining whether the sequence being revealed is fifty-fifty: e.g., one can take a (non-trivial) mixture of the fair coin measure with a measure corresponding to some other (non-extreme) bias that a coin might have.

[49]Proofs of Facts 3–5 below can be found in Belot 2013a (Sect. 4).

Fact 3. Any open-minded method for determining whether the true sequence is fifty-fifty is flummoxed by uncountably many hypotheses. (And likewise for other slippery properties.[50])

Comparisons of size need not end with questions of cardinality. We think of natural numbers as more special and rare, *qua* real numbers, than rational numbers—even though both the natural numbers and the rational numbers form countable subsets of the real numbers. Intuitively, this is because there are appreciable gaps among the natural numbers but not among the rational numbers. The topological notion of a dense subset of space explicates the intuitive notion of a subset without such gaps.[51] A subset S of our space of hypotheses is $\mathcal{H}$ is dense if and only if for each finite data set there is sequence in S that extends that data set.[52]

Fact 4. For any method M for determining whether the true sequence is fifty-fifty, the failure set of that method is dense in the space of hypotheses. (And likewise for any other particular slippery property.)

So for the problems we are interested in, open-minded methods have failure sets that are both uncountable and dense in the space of hypotheses. There is a further topological notion that allows us to recognize some uncountable dense subsets of a given space as being so small as to be essentially ignorable and others as being so large as to correspond to properties possessed by typical members of the space.[53] Thus, mathematicians will say that typical continuous functions are nowhere differentiable, while only very special functions are

[50]Note that non-open-minded methods can have countable failure sets. Consider, e.g., the problem of determining whether the true sequence encodes a binary expansion of a rational number and the method that always conjectures that it does not. This method is never flummoxed and is fooled precisely by the countable set of sequences that do encode rational numbers.

[51]Recall that a subset S of a topological space X is *dense* if every non-empty open subset of X contains a member of S, and is called *nowhere dense* if the closure of S in X contains no (non-empty) open subsets of X.

[52]Recall from fn. 7 above that for any finite data set w, B_w denotes the set of sequences that begin with w, and that we are working with a topology on $\mathcal{H}$ in which each B_w is open and in which each open set is a union of sets of this form. So a subset of $\mathcal{H}$ is dense just in case it intersects each B_w.

[53]A subset S of a topological space is *meagre* if it is a countable union of nowhere dense set; S is *residual* if its complement is meagre.

polynomials—although both the polynomials and the nowhere differentiable functions form uncountable dense subsets the space of continuous functions.[54]

Fact 5. Any open-minded method for determining whether the true sequence is fifty-fifty is flummoxed by typical hypotheses. (And likewise for any other particular slippery property.)

So far, so good—and so far, so *general.* How does all of this apply to Bayesians?

A Bayesian agent who begins with a prior P_0 defined on $\mathcal{H}$ and whose posterior after seeing n data points is P_n (the result of conditionalizing P_0 on the data seen) will conjecture at that point that the true sequence is fifty-fifty if and only if P_n assigns the set of fifty-fifty sequences a probability of at least one-half.[55] So among the methods covered by Facts 3–5 above are the methods followed by Bayesian agents.

Of course, this just puts Bayesians in the same boat with everyone else. But something sets Bayesian agents apart—a confidence problem. Bayesian agents faced with the problem of determining whether the sequence they are seeing has some particular slippery property are required to be *certain* that they will succeed: all Bayesian agents assign probability zero to the failure set of their own method when facing problems of this kind.[56]

This is odd. Suppose that God tells you and your chum Cholmondeley that you will pass your time in the Garden of Eden by trying to guess whether or not the binary sequence you are being shown is fifty-fifty (with one of you being shown a sequence of one type, the other a sequence of the other type). The problem is patently an intractable one—for open-minded methods, failure to latch on to the truth, even in the infinite long run, is the typical outcome. Here all attractive methods are bad in much the same sense in which the method of being sure come what may that the coin is fair is a bad method for determining its bias. Any normal person faced with this game will think: I need to adopt some method or other; I will look for the best one I can

[54]For a proof that the nowhere differentiable functions form a residual set, see Oxtoby 1980 (Sect. 11). The complement of this set is of course meagre—and contains the polynomials.

[55]Note that the set of fifty-fifty sequence forms a measurable subset of our space of hypotheses; see, e.g., Kechris 1995 (Sect. 11.B). From now on we restrict attention to slippery properties that correspond to measurable subsets of $\mathcal{H}$. Note that every countable subset of $\mathcal{H}$ is measurable.

[56]That is: if the agent begins with prior P_0, then the failure set of the corresponding method has P_0-measure zero. This follows, e.g., from Theorem 2 of Schervish and Seidenfeld 1990.

find; but knowing the facts about this game, I have to think that my method may well not lead me to the truth. On the Bayesian account of rationality, this line of thought is incoherent.

Contrast this situation with a couple of others in which one finds priors fully confident of success although failure is typical.

First, consider a problem in which convergence to the truth is easy: using knowledge of the outcomes of a sequence of tosses to determine whether a coin is fair, has a bias of two-thirds in favour of heads, or has a bias of two-thirds in favour of tails. No matter which hypothesis is true, any Bayesian agent who initially assigns non-zero credence to each of the three hypotheses will be more or less guaranteed to converge to the truth in this setting. Of course, no such guarantee obtains for other priors: if I initially assign probability one to the coin's being fair and update by conditionalization, then I will never assign anything other than probability zero to the other hypotheses, no matter what data I see. If I adopt such a prior, then most hypotheses have the feature that should they be true, I would never become confident in their truth—in this sense failure is typical for this prior in this setting. Nonetheless, *every* prior assigns probability zero to the set of sequences that would lead it to do anything *other* than become arbitrarily confident in the true hypothesis as more and more data is seen. Here, certain priors combine certainty of success with typical failure—and many will feel that there is something rationally defective about the priors that do so.

Consider next a problem in which *no* Bayesian agent is driven towards the truth by seeing more and more evidence: attempting to determine the denomination of a fair coin (penny, nickel, dime, quarter) from knowledge of the outcomes of a sequence of tosses. In this setting, whatever prior one begins with will be left unaltered by the data. If asked which hypothesis is most likely, one will always give the same answer, no matter what data one sees. One might happen to be correct—but of course one is typically wrong (in the sense that, most hypotheses have the feature that if they were correct, one would never consider them most plausible). The problem is a patently intractable one. Nonetheless, certain extraordinary priors are fully confident of success: if I assign probability one to the proposition that the coin is a quarter, then I assign probability zero to the set of hypotheses with the feature that, should they be true, I would fail to come to believe them to be true. Here again, it is natural to think that there is something rationally defective about priors that combine certainty of success with typical failure.

It is not, in general, a good idea to be certain of success in contexts in which failure is typical. In problems like the two just mentioned, only certain unusual priors involve this combination—priors that are, in any case, rather pathological. The situation is quite different with our slippery Property Problems—here every prior is fully confident of success, while failure is typical for open-minded priors. This represents a surprising limitation of Bayesian rationality: whereas any reasonably attractive Bayesian agent is able to appreciate that the coin denomination guessing game is an intractable one, no Bayesian agent can appreciate that the problem of guessing whether a sequence has a given slippery property is intractable. It is difficult to believe that it is rationally mandatory to consider such problems tractable.

But wait! In the problems at issue, failure is typical only for open-minded priors. Perhaps it is reasonable for other priors to be fully confident in their success? This strikes me as a desperate ploy—open-minded priors provide intuitively attractive methods for addressing these problems. Worse: there are variants of our setting in which failure is typical for *all* priors (see Belot 2013b).

But wait! For all that has been said here, it is not clear that the notion of typicality at issue in Fact 5 should be of interest in this context. That is true. But there is more to be said by way of motivating this notion (see Belot 2013a). And there is a long tradition among Bayesian statisticians of looking askance at priors which fail to converge to the truth for hypotheses that are typical in this sense.[57]

But wait! At best the above shows that orthodox subjective Bayesianism and its tempered alternatives have a problem—what about generalizations of the orthodox framework in which rational agents are allowed to have credences representable by objects more general than probability measures? Here the situation is less clear: the objection developed above turns on the availability of theorems (such as the result cited in fn. 56 above) that show that each Bayesian agent is subjectively certain of convergence to the truth in the infinite long run. But most work on asymptotic questions has been conducted by statisticians working in the standard framework and thus concerns probability measures.[58]

[57]See, e.g., the classic papers Freedman 1963 and 1965 and Diaconis and Freedman 1986.

[58]For discussion and references concerning the case of merely finitely additive probability measures, see Zabell 2002. Regarding primitive conditional probabilities: note that Gaifman (1986, p. 341) reports that the paradigmatic convergence result of Gaifman and Snir 1982 carries over unchanged to this setting.

But there are some intriguing positive results. Weatherson (forthcoming) shows how to construct an imprecise prior (=a set of probability measures) that does not consider itself certain to succeed prior to inquiry when facing a version of the problem discussed above. Elga (manuscript) does the same in the setting of merely finitely additive probability measures. This is encouraging. But it is far from a sufficient response to the problem for anyone interested in defending some form of generalized Bayesianism as an account of rationality, as opposed to as a theory of statistical inference. The point of the contrast is this: a Bayesian theory of statistical inference just needs to tell you how to behave in the face of data within a given problem context—one looks at the problem being asked and then selects a prior suitable for one's purposes. A theory of rationality aims to do more: if God tells you in the Garden of Eden that you are to be shown a binary sequence, one bit at a time, then at that point there must be some sort of (generalized) prior that sums up your credences about which sequence you will be shown—and you are rationally committed to forming beliefs by conditionalizing on that prior. You do not get to use one prior if God goes on to ask you after each bit is revealed how likely it is that the sequence being revealed encodes a rational number and another prior if you are instead asked how likely it is that the sequence is fifty-fifty. Those seeking to defend one or another form of generalized Bayesian account of rationality here have their work cut out for them.

6. Difficulty Itself

We started with the observation that there is a tension between objectivity and bias, and with the idea that it would be nice to identify those forms of objectivity that are obtainable and to understand the range of biases compatible with them. We saw that strict and long-run objectivity are beyond our reach: there is no way to conduct genuine inquiry in the absence of bias; and it is too much to ask of a method that it should always eventually be able to settle on the truth. This led us to an investigation of asymptotic objectivity. Can we find methods that lead to conjectures that are more or less guaranteed to converge to the truth as more and more data points are seen?

For some problems, this can easily be achieved. We might be interested in determining the propensity of a coin to come up heads and expect to see a sequence of outcomes of tosses. Or we might be trying to guess the identity of a binary sequence that is being revealed to us one bit at a time. For problems like this it is easy to come up with asymptotically objective methods—the problem is that there are far too many of them, that it is difficult to choose between them, and that they differ all too often in their recommendations at finite stages of inquiry.

But we have seen that even when asymptotically objective methods exist for identifying fine-grained hypotheses, there may be harder problems lurking: for instance, there can be no asymptotically objective method for determining whether the binary sequence being revealed to you is fifty-fifty (i.e., split evenly between zeroes and ones in the infinite long run). And any method that is open-minded in a certain sense (never making up its mind irredeemably about this problem on the basis of a finite amount of evidence) is flummoxed by typical sequences in the space of hypotheses, flip-flopping *ad infinitum* between YES and NO.

This phenomenon causes trouble. For those who would follow Reichenbach in taking asymptotic objectivity to be the characteristic virtue of good methods, a vicious closure problem follows—they have to maintain, for example, that we can be warranted in believing that the true sequence is 0101010... (or be willing to act as if we do), but are forbidden from taking the same attitude towards the claim that zero and one occur with the same frequency in the sequence. If we want anything like our own methods to count as objective, we need to allow that asymptotic objectivity is not a necessary condition for objectivity.

The same phenomenon causes problems for other approaches. At the opposite end of the spectrum from Reichenbachians, who often find themselves embarrassed by a surfeit of methods approved by their official standard, lie uniquarians, who maintain that there is only one rational doxastic response to any given body of evidence. If we identify the evidence available to a curve-fitter with the data points seen, then this leads us to expect that the unique rational method for determining whether the sequence being revealed is fifty-fifty must be optimal with respect to its domain of asymptotic objectivity. But it is demonstrable that optimality is unobtainable here—there is better and better but no best. So uniquarians must deny, on pain of incoherence, that the evidence available to a curve-fitter should be identified with the data points seen. But to go down that road is to render epistemology a curiously limited enterprise, with nothing to say about how one should respond to the sort of data sets and questions prevalent in the sciences.

Somewhere between the relatively wide range of biases permitted by Reichenbachians and very narrow range imagined by uniquarians, lies the realm of subjective Bayesians, who build all degrees of freedom into the choice of an initial probability distribution over the relevant space of hypotheses. Bayesian agents run into a special difficulty when confronting problems like that of determining whether a binary sequence is fifty-fifty—they are sure that they will succeed, even though failure is typical for this task. It is difficult to believe that this sort of self-confidence can be rationally mandatory. (And it is note-worthy that when facing other sorts of intractable problems, reasonable Bayesian agents are able to arrive at reasonable views about their chances of success.)

So what sort of account of obtainable objectivity and acceptable bias should we be looking for? One that allows for objective inquiry in the absence of a guarantee of convergence to the truth, under at least some conditions. One that neither postulates a single standard for all possible beings, nor severs all connections between epistemology, on the one hand, and the sciences and philosophy of science, on the other. One that allows us to take a realistic view of our chances of success when faced with intractable problems.

Appendix

Above we considered agents who seek to learn facts about a binary sequence that is being revealed to them one bit at a time. This constitutes the simplest possible model of empirical inquiry in the regime in which the amount of in-principle relevant evidence is inexhaustible. In this sense, it provides a natural idealization of scientific investigation. At the same time, it is natural to worry that the difficulties encountered above might be artifacts of certain special features of the problems considered above. In physics, for example, data typically come in the form of real or complex numbers and the spaces of hypotheses in play are usually spaces of of real- or complex-valued functions—spaces that are quite different in structure from the space $\mathcal{H}$ of binary sequences. Further, in the discussion above, we relied heavily on the fact that $\mathcal{H}$ includes *all* binary sequences, no matter how well- or ill-behaved. But in scientific contexts, one almost always restricts attention to functions that are well-behaved in one or another sense (e.g., by requiring certain degrees of smoothness or by imposing boundary conditions). So it is natural to wonder whether the difficulties we have been concerned with arise in more sophisticated settings. We will see that they do.

First some background. Consider the space $C^\infty(S^1)$ of infinitely-differentiable, real-valued functions on the circle (we will think of these as defined on the closed interval $[-\pi, \pi]$ subject to the condition that they take the same value at each endpoint).[59] Recall that functions

[59] The same argument would work for $C^\infty(T^n)$, where T^n is an n-dimensional torus.

in $C^\infty(S^1)$ have well-defined Fourier expansions. For $f \in C^\infty(S^1)$ we define

$$a_k := \frac{1}{\pi}\int_{-\pi}^{\pi} f(\theta)\cos k\theta \; d\theta \qquad (k = 0, 1, 2, \ldots)$$
$$b_k := \frac{1}{\pi}\int_{-\pi}^{\pi} f(\theta)\sin k\theta \; d\theta \qquad (k = 1, 2, 3, \ldots).$$

and call a_0, (a_1, b_1), (a_2, b_2), and so on, the *Fourier coefficients* of f. We call

$$\frac{a_0}{2} + \sum_{k=1}^{k=\infty} (a_k \cos kx + b_k \sin kx)$$

the *Fourier series* of f. In the context of Fourier analysis, it is natural to equip $C^\infty(S^1)$ with the norm

$$||f|| = (|a_0|^2 + \sum |a_k|^2 + |b_k|^2)^{\frac{1}{2}}$$

(the norm of a function in $C^\infty(S^1)$ is a non-negative real number). This norm induces a notion of convergence for sequences of functions in $C^\infty(S^1)$: $f_n \to f$ if and only if $||f - f_n|| \to 0$ as $n \to \infty$. Note, in particular, that this provides one of the several senses in which the Fourier series of a function $f \in C^\infty(S^1)$ converges to f: let $S_N(x)$ denote $\frac{a_0}{2} + \sum_{k=1}^{k=N}(a_k \cos kx + b_k \sin kx)$, the Nth *partial sum* of the Fourier series of f; then $||f - S_N|| \to 0$ as $N \to \infty$.

Now we can set up our problems. A curve-fitter is gradually shown the Fourier coefficients of some function in $C^\infty(S^1)$ (first a_0, then a_1 and b_1, then a_2 and b_2, and so on). At each stage the agent is asked to conjecture the true identity of the function generating this data and also to guess whether this function has a finite Fourier expansion (i.e., whether all but finitely many of its Fourier coefficients vanish). This new model, while still very simple, features a space of hypotheses that is both rich enough and sufficiently constrained to seem like a decent stand-in for the sorts of spaces of hypotheses that arise in scientific practice. And the sort of data involved is also a good stand-in for the sorts of data that arise (for instance) in geophysics (see Parker 1994, Ch. 2).

We find something familiar. It is easy to come up with an asymptotically reliable method for identifying the function f generating the data. For example, one can at each stage just guess the function whose only non-zero Fourier coefficients coincide with the data one has been shown so far—this is asymptotically reliable because, as noted above, the partial sums S_N of the Fourier series of any function in $C^\infty(S^1)$ converge to that function.

Fact 6. There can be no asymptotically reliable method for determining whether the function generating the data has a finite Fourier expansion.[60]

Proof. The argument is parallel to that for Fact 1 above. Suppose that M is an asymptotically reliable method for determining whether the function generating the data has a finite Fourier expansion.

(i) No matter which function is generating the data, M will have to settle down after a finite time to the correct answer.
(ii) Any finite data set is consistent with either answer.
(iii) So if after seeing a given data set, M conjectures that the true function does not have a finite Fourier expansion, it is always possible to get M to change its mind by showing it a sufficiently large block of data consisting of nothing but zeroes; likewise, whenever M conjectures that the true function does have a finite Fourier expansion, it is always possible to get it to change its mind by showing it a sufficiently large block of non-vanishing coefficients.
(iv) Here is a recipe for constructing a data stream that will cause M to flip-flop *ad infinitum.* Let $f \in C^\infty(S^1)$ be a function with Fourier coefficients a_0, a_1, b_1, ..., all of which are non-zero.We define a sequence a_0^*, a_1^*, b_1^*, ... as follows:
 a) $a_0^* = a_0$;
 b) $(a_{k+1}^*, b_{k+1}^*) = (a_{k+1}, b_{k+1})$ $(k \geq 0)$ if, when shown a_0^*, ..., b_k^*, M guesses that the true function has a finite Fourier expansion;
 c) $(a_{k+1}^*, b_{k+1}^*) = (0, 0)$ otherwise.
 That is, we construct the sequence a_0^*, a_1^*, b_1^*, ... by alternating blocks of data consisting of the Fourier coefficients of f with blocks of data consisting of all zeroes in such a way as to ensure that M flip-flops *ad infinitum.*

[60]This lets us make good on a promise made in Sect. 1 above. In the present model, the space of hypotheses and the space of finite data sets both have the cardinality of the continuum. But there can be no method for guessing the true identity of the function generating the data that is long-run objective in the sense of Sect. 1 above. For if there were, that method would allow us to construct an asymptotically reliable method for determining whether the true function has a finite Fourier expansion (conjecture that it does if and only if the conjecture of the long-run objective method for the Identification Problem has this property)—but there can be no such method. So here we have a context in which long-run objectivity is unobtainable even though there is no mismatch in cardinality between the space of hypotheses and the space of finite data sets.

(v) The a_0^*, a_1^*, b_1^*, ... are the Fourier coefficients of a function f^* in our space of hypotheses $C^\infty(S^1)$. This is because a sequence $(a_0, a_1, b_1, \ldots)$ gives the Fourier coefficients of a function in $f \in C^\infty(S^1)$ if and only if the a_k and b_k go to zero sufficiently quickly (see, e.g., Taylor 2011, Sect. 3.1)—and that is a feature that cannot be spoiled by setting some of the elements of the sequence to zero.

(vi) So there is a function on our space of hypotheses that generates a data stream that flummoxes M.

The contradicts our original assumption that M was asymptotically reliable for this problem. So there can be no such M. □

This shows that even if we restrict attention to 'nice' hypotheses, we still get a counterpart of Fact 1.[61] And a counterpart of Fact 2 also holds (the argument given above for Fact 2 carries over directly). It follows that every method for trying to determine whether the function generating the data has a finite Fourier expansion will have an infinite failure set. Further, there is no obstacle to applying the Bayesian framework to this problem (the set of functions with finite Fourier expansions is a Borel subset of $C^\infty(S^1)$ and hence measurable).

Weaker versions of Facts 3 and 4 hold in the present setting. Call a method *regular* if, no matter what data it has seen so far, it can always be brought to guess that the function f generating the data has a finite Fourier expansion by being shown a sufficiently large block of vanishing Fourier coefficients and can always be brought to guess that f doesn't have a finite Fourier expansion by being shown a sufficiently large block of non-vanishing Fourier coefficients.

Fact 7. The flummoxers of any regular method M form an uncountable dense subset of $C^\infty(S^1)$.

Proof. First note that for any $f \in C^\infty(S^1)$ and any $\varepsilon > 0$, there is an N such that $||f - S_N|| < \varepsilon$. And if $g \in C^\infty(S^1)$ is a function whose Fourier series is identical to that of f except maybe that some of the coefficients with order higher than N have been set to zero, then $||f - g|| < \varepsilon$ as well. This means that for any regular method M and any $f \in C^\infty(S^1)$ we can construct a sequence $(f_1, f_2, \ldots)$ of flummoxers of M that coverage to f as follows: choose each f_k to match f in enough of its initial

[61]The same sort of argument works for other classes of functions admitting Fourier or other sorts of expansions, so long as the class of functions can be characterized as those whose expansion coefficients fall off at least as fast as some target rate. For instance, we could work certain Sobolev spaces and with Fourier expansions (see Taylor 2011, Sect. 4.3) or with Schwartz class functions and Hermite expansions (see Simon 1971, Theorem 1).

Fourier coefficients to ensure that $||f - f_k|| < \frac{1}{k}$, given that the remaining Fourier coefficients are going to be chosen as in the proof of Fact 6 above to ensure that f_k flummoxes M. So the flummoxers of M are dense in $C^\infty(S^1)$.

Next, let α be a real number between zero and one with binary expansion $.\alpha_1\alpha_2\alpha_3 \ldots$. Let M be a regular method and let f be a function in $C^\infty(S^1)$ all of whose Fourier coefficients are non-zero. We construct a flummoxer for M as follows: show M the first Fourier coefficient of f; look at what M conjectures; then show M a block of coefficients that makes it change its mind (either a block of zeroes or a block of Fourier coefficients of f); the first coefficient it sees after it changes its mind is α_1 times the next coefficient of f; followed by a (possibly empty) set of coefficients designed to make it change its mind a second time; then α_2 times the next coefficient of f; and so on. In this way we construct a distinct flummoxer for M for every real number between zero and one. $\square$

It is natural to hope that Fact 5 also carries over in some form.

References

Anderson, Elizabeth 2004: 'Uses of Value Judgments in Feminist Social Science: A General Argument, with Lessons from a Case Study of Feminist Research on Divorce'. *Hypatia,* 19, pp. 1–24.

Antony, Louise (ed) 2007: *Philosophers Without Gods: Meditations on Atheism and the Secular.* Oxford: Oxford University Press.

Belot, Gordon 2013a: 'Bayesian Orgulity'. *Philosophy of Science,* 80, pp. 483–503.

Belot, Gordon 2013b: 'Failure of Calibration is Typical'. *Statistics and Probability Letters,* 83, pp. 2316–18.

Belot, Gordon (forthcoming): 'Curve-Fitting for Bayesians?' Forthcoming in *British Journal for the Philosophy of Science.*

Boghossian, Paul 2008: 'Replies to Wright, MacFarlane, and Sosa'. *Philosophical Studies,* 141, pp. 409–32.

Braithwaite, Richard (ed.) 1931: *Foundations of Mathematics and Other Essays.* London: Kegan, Paul, Trench, Trubner, & Co.

Carnap, Rudolf 1945: 'On Inductive Logic'. *Philosophy of Science,* 12, pp. 72–97.

Carnap Rudolf 1950: 'Empiricism, Semantics, and Ontology'. *Revue Internationale de Philosophie,* 4, pp. 20–40.

Carnap, Rudolf 1962: 'The Aims of Inductive Logic'. In Nagel, Suppes, and Tarski 1962, pp. 303–18.

Chalmers Alan 1976: *What is This Thing Called Science?* Indiana: Hackett, third edn 1999.

Cohen, Robert S. and Larry Laudan (eds) 1983: *Physics, Philosophy, and Psychoanalysis: Essays in Honor of Adolf Grünbaum.* Dordrecht: Reidel.

Conant, James and John Haugeland (eds) 2000: *Thomas S. Kuhn: The Road Since Structure.* Chicago: University of Chicago Press.

Conee, Earl and Richard Feldman 2004: *Evidentialism: Essays in Epistemology.* Oxford: Oxford University Press.

Conee, Earl and Richard Feldman 2008: 'Evidence'. In Smith 2008, pp. 83–104.

Daston, Lorraine and Peter Galison 2007: *Objectivity.* Cambridge, MA: Zone Books.

Diaconis, Persi and David Freedman 1986: 'On the Consistency of Bayes Estimates'. *The Annals of Statistics,* 14, pp. 1–26.

Dretske, Fred 2005a: 'The Case Against Closure'. In Steup and Sosa 2005, pp. 13–26.

Dretske, Fred 2005b: 'Reply to Hawthorne'. In Steup and Sosa 2005, pp. 43–6.

Earman, John 1992: *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory.* Cambridge, MA: MIT Press.

Earman, John 1993: 'Underdetermination, Realism, and Reason'. *Midwest Studies in Philosophy,* XVIII, pp. 19–38.

Elga, Adam (manuscript): 'Bayesian Humility'. Unpublished.

Enoch, David and Joshua Schechter 2008: 'How Are Basic Belief-Forming Methods Justified?' *Philosophy and Phenomenological Research,* 76, pp. 547–79.

Feigl, Herbert 1934: 'The Logical Character of the Principle of Induction'. *Philosophy of Science,* 1, pp. 20–9.

Feigl, Herbert 1954: 'Scientific Method without Metaphysical Presuppositions'. *Philosophical Studies,* V, pp. 17–29.

Feigl, Herbert and May Brodbeck (eds) 1953: *Readings in the Philosophy of Science.* New York: Appleton–Century–Crofts.

Feldman, Richard 2007: ‘Reasonable Religious Disagreements.’ In Antony, 2007, pp. 194–214.

Feldman, Richard and Earl Conee 2001: ‘Internalism Defended’, in Conee and Feldman 2004, pp. 53–82. Originally published in *American Philosophical Quarterly,* 38.

Fetzer, James (ed) 2001: *The Philosophy of Carl G. Hempel: Studies in Science, Explanation, and Rationality.* Oxford: Oxford University Press.

Field, Hartry 2009: ‘The Normative Role of Logic’. *Aristotelian Society Supplementary Volume,* 83, pp. 251–68.

Forster, Malcolm and Elliott Sober 1994: ‘How to Tell when Simpler, More Unified, or Less *ad hoc* Theories will Provide More Accurate Predictions’. *British Journal for the Philosophy of Science,* 45, pp. 1–35.

Freedman, David 1963: ‘On the Asymptotic Behavior of Bayes’ Estimates in the Discrete Case’. *Annals of Mathematical Statistics,* 34, pp. 1386–403.

Freedman, David 1965: ‘On the Asymptotic Behavior of Bayes’ Estimates in the Discrete Case. II’. *Annals of Mathematical Statistics,* 36, pp. 454–6.

Gabbay, Dov, Stephan Hartmann, and John Woods (eds) 2010: *Inductive Logic.* Amsterdam: Elsevier.

Gaifman, Haim 1986: ‘Towards a Unified Concept of Probability’. In Marcus, Dorn, and Weingartner 1986, pp. 319–50.

Gaifman, Haim and Marc Snir 1982: ‘Probabilities Over Rich Languages’. *Journal of Symbolic Logic,* 47, pp. 495–548.

Glymour, Clark 1980: *Theory and Evidence.* Princeton: Princeton University Press.

Hacking, Ian 1968: ‘One Problem about Induction’. In Lakatos 1968, pp. 44–59.

Hacking, Ian 2000: ‘How Inevitable are the Results of Successful Science?’ *Philosophy of Science,* 67, pp. S58–71.

Hájek, Alan 2008: ‘Arguments for—or against—Probabilism?’ *British Journal for the Philosophy of Science,* 59, pp. 793–819.

Hempel, Carl Gustav 1966: *Philosophy of Natural Science.* Upper Saddle River, NJ: Prentice–Hall.

Hempel, Carl Gustav 1983: 'Valuation and Objectivity in Science'. In Fetzer 2001, pp. 372–95. Originally published in Cohen and Laudan 1983.

Hookway, Christopher 2007: 'Fallibalism and the Aim of Inquiry'. *Aristotelian Society Supplementary Volume,* 81, pp. 1–22.

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

Howson, Colin and Paul Urbach 1989: *Scientific Reasoning: The Bayesian Approach* (first edition). Chicago: Open Court.

Howson, Colin and Paul Urbach 2006: *Scientific Reasoning: The Bayesian Approach* (third edition). Chicago Open Court.

James, William 1902: *The Varieties of Religious Experience.* London: Longmans, Green, & Co.

Jeffreys, Harold 1933: 'Probability, Statistics, and the Theory of Errors'. *Proceedings of the Royal Society of London,* 140, pp. 523–35.

Joyce, James 2004: 'Bayesianism'. In Mele and Rawling, 2004, pp. 132–155.

Joyce, James 2010: 'The Development of Subjective Bayesianism'. In Gabbay, Hartmann, and Woods 2010, pp. 415–76.

Kechris, Alexander 1995: *Classical Descriptive Set Theory.* Berlin: Springer–Verlag.

Kelly, Kevin 1996: *The Logic of Reliable Inquiry.* Oxford: Oxford University Press.

Kelly, Kevin and Clark Glymour 1989: 'Convergence to the Truth and Nothing but the Truth'. *Philosophy of Science,* 56, pp. 185–220.

Kemeny, John 1953: 'The Use of Simplicity in Induction.' *Philosophical Review,* 62, pp. 391–408.

Keynes, John Maynard 1921: *A Treatise of Probability.* London: MacMillan.

Kuhn, Thomas (1970) 'Reflections on My Critics'. In Conant and Haugeland 2000, pp. 123–75. Originally published in Lakatos and Musgrave 1970.

Kuhn, Thomas (1963) *The Structure of Scientific Revolutions.* Chicago: University of Chicago Press, second edn 1970.

Lakatos, Imre (ed.) 1968: *The Problem of Inductive Logic.* Amsterdam: North–Holland.

Lakatos, Imre and Alan Musgrave (eds) 1970: *Criticism and the Growth of Knowledge.* Cambridge: University of Cambridge Press.

Leslie, Sarah-Jane 2013: 'Essence and Natural Kinds: When Science Meets Preschooler Intuition.' *Oxford Studies in Epistemology,* 4, pp. 108–65.

Madden, Edward 1964: 'Peirce on Probability' In Moore and Robin 1964, pp. 122–40.

Maher, Patrick 2006: 'The Concept of Inductive Logic'. *Erkenntnis,* 65, pp. 185–206.

Marcus, Ruth Barcan, George Dorn, and Paul Weingartner (eds) 1986: *Logic, Methodology and Philosophy of Science VII.* Amsterdam: Elsevier.

Mele, Alfred and Piers Rawling (eds) 2004: *The Oxford Handbook of Rationality.* Oxford: Oxford University Press.

Moore, Edward and Richard Robin (eds) 1964: *Studies in the Philosophy of Charles Sanders Peirce* (second series). Amherst: University of Massachusetts Press.

Moss, Sarah (forthcoming): 'Credal Dilemmas'. Forthcoming in *Noûs.*

Nagel, Ernest, Patrick Suppes, and Alfred Tarski (eds) 1962: *Logic, Methodology, and Philosophy of Science: Proceedings of the 1960 International Congress.* Stanford: Stanford University Press.

Oxtoby, John 1980: *Measure and Category: A Survey of the Analogies between Topological and Measure Spaces.* Berlin: Springer–Verlag.

Parker, Robert 1994: *Geophysical Inverse Theory.* Princeton: Princeton University Press.

Putnam, Hilary 1963a: '"Degree of Confirmation" and Inductive Logic'. In his 1979, 270–92. Originally published in Schilpp 1963.

Putnam, Hilary (1963b) 'Probability and Confirmation'. In his 1979, pp. 293–304. Originally Published as *The Voice of America, Forum Philosophy of Science 10.* Washington: U.S. Information Agency.

Putnam, Hilary 1979: *Mathematics, Matter and Method.* Cambridge: Cambridge University Press.

Railton, Peter 1994: 'Truth, Reason, and the Regulation of Belief'. *Philosophical Issues,* 5, pp. 71–93.

Ramsey, Frank 1931: 'Truth and Probability.' In Braithwaite 1931, 156—98.

Reichenbach, Hans 1933: 'The Logical Foundations of the Concept or Probability'. In Feigl and Brodbeck 1953, pp. 456–74. Originally published as 'Die logischen Grundlagen des Wahrscheinlichkeitsbegriffs,' *Erkenntnis,* 3.

Reichenbach, Hans 1938: *Experience and Prediction.* Chicago: University of Chicago Press.

Reichenbach, Hans 1939: 'Dewey's Theory of Science'. In Schilpp 1939, pp. 157–92.

Reichenbach, Hans 1949: *Theory of Probability.* Berkeley: University of California Press.

Rosen, Gideon 2001: 'Nominalism, Naturalism, Epistemic Relativism'. *Philosophical Perspectives,* 15, pp. 69–91.

Salmon, Wesley 1991: 'Hans Reichenbach's Vindication of Induction'. *Erkenntnis,* 35, pp. 99–122.

Salmon, Wesley 2005a: 'Rationality and Objectivity in Science.' In his 2005b, pp. 93–116.

Salmon, Wesley 2005b: *Reality and Rationality.* Oxford: Oxford University Press.

Savage, Leonard 1972: *Foundations of Statistics.* London: Methuen.

Schaffer, Jonathan 2005: 'Constrastive Knowledge'. *Oxford Studies in Epistemology,* 1, pp. 235–71.

Schervish, Mark and Teddy Seidenfeld 1990: 'An Approach to Consensus and Certainty with Increasing Evidence'. *Journal of Statistical Planning and Inference,* 25, 401–14.

Schilpp, Paul (ed) 1939: *The Philosophy of John Dewey.* Evanston: Northwestern University Press.

Schilpp, Paul (ed) 1963: *The Philosophy of Rudolf Carnap.* Chicago: Open Court.

Schoenfeld, Miriam 2014: 'Permission to Believe: Why Permissivism is True and What It Tells Us About Irrelevant Influences on Belief'. *Noûs,* 48, 193–218.

Simon, Barry 1971: 'Distributions and Their Hermite Expansions'. *Journal of Mathematical Physics,* 12, pp. 140–8.

Skyrms, Brian 1991: 'Carnapian Inductive Logic for Markov Chains'. *Erkenntnis,* 35, pp. 439–60.

Smith, Quentin (ed) 2008: *Epistemology: New Essays.* Oxford: Oxford University Press.

Sober, Elliott 1988: 'Likelihood and Convergence'. *Philosophy of Science,* 55, pp. 228–37.

Steup, Matthias and Ernest Sosa (eds.) 2005: *Contemporary Debates in Epistemology.* Oxford: Blackwell.

Taylor, Michael 2011: *Partial Differential Equations,* Volume I. Berlin: Springer–Verlag.

Weatherson, Brian (forthcoming) 'For Bayesians, Rational Modesty Requires Imprecision'. Forthcoming in *Ergo.*

Weinberg, Steven 2001: *Facing Up: Science and Its Cultural Adversaries.* Cambridge, MA: Harvard University Press.

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

White, Roger 2009: 'Evidential Symmetry and Mushy Credence.' *Oxford Studies in Epistemology,* 3, pp. 161–86.

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

Williamson, Timothy 2008: 'Why Epistemology Cannot be Operationalized'. In Smith 2008, pp. 277–300.

Wright, Crispin 2004: 'Warrant for Nothing (and Foundations for Free)?' *Aristotelian Society Supplementary Volume,* 78, pp. 167–212.

Zabell, Sandy 1989: 'The Rule of Succession'. *Erkenntnis,* 31, pp. 283–321.

Zabell, Sandy 2002: 'It All Adds Up: The Dynamic Coherence of Radical Probabilism'. *Philosophy of Science,* 69, pp. S98–103.

Department of Philosophy, University of Michigan, belot@umich.edu