

# There Is No Pure Empirical Reasoning

## 1. Empiricism and the Question of Empirical Reasons

Empiricism may be defined as the view there is no a priori justification for any synthetic claim. Critics object that empiricism cannot account for all the kinds of knowledge we seem to possess, such as moral knowledge, metaphysical knowledge, mathematical knowledge, and modal knowledge.<sup>1</sup> In some cases, empiricists try to account for these types of knowledge; in other cases, they shrug off the objections, happily concluding, for example, that there is no moral knowledge, or that there is no metaphysical knowledge.<sup>2</sup>

But empiricism cannot shrug off just any type of knowledge; to be minimally plausible, empiricism must, for example, at least be able to account for paradigm instances of *empirical* knowledge, including especially scientific knowledge.

Empirical knowledge can be divided into three categories: (a) knowledge by direct observation; (b) knowledge that is deductively inferred from observations; and (c) knowledge that is non-deductively inferred from observations, including knowledge arrived at by induction and inference to the best explanation. Category (c) includes all scientific knowledge. This category is of particular import to empiricists, many of whom take scientific knowledge as a sort of paradigm for knowledge in general; indeed, this forms a central source of motivation for empiricism.<sup>3</sup> Thus, if there is any kind of knowledge that empiricists need to be able to account for, it is knowledge of type (c).

I use the term “empirical reasoning” to refer to the reasoning involved in acquiring this type of knowledge – that is, to any instance of reasoning in which (i) the premises are justified directly by observation, (ii) the reasoning is non-deductive, and (iii) the reasoning provides adequate justification for the conclusion. I assume that non-deductive reasoning justifies its conclusion, if at all, by rendering the conclusion probable. I call a case of empirical reasoning “pure” if it also satisfies the further condition (iv) that the reasoning does not depend upon or in any other way imply the existence of any a priori justification for any synthetic claim.

I contend that there is no pure empirical reasoning. There are cases of reasoning satisfying (i)-(iii), but none of them also satisfy (iv). All empirical

---

<sup>1</sup> Bonjour 1998; Bealer 1992; Huemer 2005, ch. 5.

<sup>2</sup> Carnap 1932; Ayer 1952, chs. 1, 6.

<sup>3</sup> Think, for example, of the seemingly unconscious way in which Quine (1951, p. 39) conflates “the totality of our so-called knowledge or beliefs” with “total science” or “the whole of science,” and Schlick (1974, p. ix) conflates “the entire system of knowledge” with “all the sciences.” Schlick (p. x) immediately goes on to claim that work in epistemology should rest on examining natural science as a model.

reasoning requires background cognitive attitudes that cannot be justified on the basis of observations, logic, or the analysis of concepts. (Non-skeptical) empiricism is thus an untenable epistemological doctrine.

## 2. Empiricism Has No Coherent Account of Empirical Reasons

### 2.1. Empirical Reasoning Requires Background Probabilities

Any empirical inference depends for its justificatory force on background, probabilistic information. This information is of two kinds: (i) information reflected by the prior probability of the conclusion; (ii) information reflected in certain conditional probabilities, or in prior probabilities of certain conjunctive propositions (these last two being interchangeable in probability theory).<sup>4</sup> I start with the theoretical justification of this point, followed by illustrative examples.

The theoretical justification begins from Bayes' Theorem:

$$P(b | e) = \frac{P(b) \cdot P(e | b)}{P(b) \cdot P(e | b) + P(\sim b) \cdot P(e | \sim b)} \quad \text{Equation 1}$$

Let  $e$  be some observational evidence and  $b$  be the conclusion of some empirical reasoning (perhaps inductive or abductive) from that evidence. The reasoning succeeds in justifying  $b$  only if  $b$  is rendered probable by  $e$ , which is true only if  $P(b | e)$  is sufficiently high (and higher than  $P(b)$ ). And that, according to Equation 1, depends on the values of  $P(b)$ ,  $P(e | b)$ , and  $P(e | \sim b)$ . Some combinations of values for those three probabilities let  $b$  be highly probable in the light of  $e$ , whereas others do not. Therefore, the justificatory force of the empirical reasoning from  $e$  to  $b$  depends on there being (justification for) a suitable probability distribution.

This reasoning turns on no unduly strong assumptions. In particular, note

- i. The reasoning does not depend on the assumption that epistemic justification can be reduced to probability. It requires at most the assumption that, for a conclusion to be justified via empirical reasoning, it is a *necessary condition* that the conclusion be probable (or at least, not highly improbable), given the premises.
- ii. Nor does the argument assume that the mathematical principles of probability tell us all we need to know about justification, or even about probability. The argument requires only that the theorems of probability,

---

<sup>4</sup> That is, modulo some set of prior probabilities for atomic propositions, if one is given a complete set of conditional probabilities, one can compute a complete set of probabilities of conjunctions, and vice versa. So in empirical reasoning, one may rely either on credences for conjunctions or on conditional credences.

especially Bayes' Theorem, are true.

- iii. Nor need we assume that, in empirical reasoning, subjects must reason about probabilities, hold beliefs about probabilities, or be able to assign precise numerical values to the relevant probabilities. The claim is only that the reasoner must have available some justification (independent of the conclusion) that would suffice to reject probability distributions on which the conclusion of the inference is improbable in light of the evidence.

For an illustration of the dependence of empirical reasoning on prior probabilities, compare two hypotheses that might explain my sensory experiences:

- a. The brain-in-a-vat hypothesis (BIVH): I am a brain in a vat who is being stimulated by the scientists in such a way as to create a perfect simulation of life in the early twenty-first century.
- b. The real-world hypothesis (RWH): I am a normal person living in the early twenty-first century.

These hypotheses do not differ in their empirical predictions.<sup>5</sup> Nevertheless, it is reasonable to believe RWH and reject BIVH. Indeed, if we met someone who believed BIVH, we would consider that person irrational to the point of insanity. But since the two theories are empirically equivalent, the reason why one of them is rational and the other not cannot lie in differing empirical support. It must therefore be *a priori*. When non-philosophers hear of BIVH, a common reaction is that it is a “crazy” scenario. This reaction can be naturally interpreted as communicating that one attaches to the BIV scenario a very low initial credence.

For those who dislike flights of philosophical fancy, here is a more scientific example. In 2011, the noted Cornell psychologist Daryl Bem published a paper reporting statistical evidence for the existence of psychic phenomena – specifically, a form of precognition involving backwards causation.<sup>6</sup> The precognition hypothesis had passed tests of statistical significance in eight out of nine experiments. For example, in one experiment, subjects were to guess which of two curtains would have a picture behind it, where the picture and its location were randomly selected by a computer *after* the subject made their guess. Bem found that, in the cases where an erotic image was going to be

---

<sup>5</sup> One might claim that the two hypotheses differ because RWH predicts that I will *observe* physical objects (where “observe” is a success term), while BIVH predicts that I won't genuinely observe anything. The success of this move turns on what we count as our foundational empirical evidence, whether physical-object propositions or propositions about experiences, internally construed. If you are tempted by this move, substitute another example of two empirically equivalent theories, where one of the theories nevertheless strikes us as crazy – for example, the sodium-uranium transmutation theory discussed in section 3.4 below.

<sup>6</sup> Bem 2011.

displayed, subjects did significantly better than chance at guessing where the picture would appear – suggesting, according to Bem, that the attractiveness of the picture that the subject would later see influenced the subject’s guess.

When I first heard these results, I was skeptical of the theory. More than skeptical, in fact: I did not provisionally accept the theory pending further investigation, nor did I withhold judgment. Rather, I thought the theory was false, and that is what I still think. In essence, I believe that Bem’s experimental results were due to chance, rather than to precognition. Why? Prior to this experiment, I find the existence of backwards causation and precognition extremely improbable; thus, it is more plausible that Bem obtained statistically significant results by chance.<sup>7</sup>

This sort of reaction is common in response to alleged evidence for psychic phenomena: skeptics often cite “coincidence” as a hypothesis superior to the postulation of psychic phenomena. But we would not react in this way if some more initially plausible hypothesis were at issue. Extraordinary claims, as we say, require extraordinary evidence. The lesson is that the success of an empirical argument depends on the prior probability of the conclusion.

Now to illustrate the dependence on *conditional* probabilities, assume that I have observed 100 shamrocks (which I can identify independent of their color), all of which I found to be green. I might infer that, since all the observed shamrocks so far have been green, the next one will also be green.

But now consider a new color predicate, “grue”, where it is stipulated that an object counts as grue just in case: it is observed before now and is green, or it is not observed before now and is blue.<sup>8</sup> It is true that every shamrock I have observed hitherto has been green, but it is also true that they have all been *grue*. So I might infer that the next one will be grue as well.

Formally, the two inferences are analogous. But the rational conclusion is that the next shamrock will be green, not grue. What is the difference between the “green” hypothesis and the “grue” hypothesis?

In this case, the difference lies either in conditional probabilities, or in probabilities of certain conjunctions (these being mathematically interchangeable here): the conditional probability of unobserved shamrocks being *grue* given that the observed shamrocks were grue is much lower than the conditional probability of unobserved shamrocks being *green* given that the observed shamrocks were green. Alternately, we may say: the prior probability of both observed and unobserved shamrocks being uniform with respect to greenness is higher than the prior probability of their being uniform with respect to grueness. Presumably, this contrasting treatment of “grue” and “green” hypotheses in our probability distributions reflects our tendency to think such things as that green is a genuine, natural property in a sense in which grue is not, that the predicate “green” cuts nature at its joints in a way that

---

<sup>7</sup> For a more sophisticated version of this view, see Rouder and Morey 2011.

<sup>8</sup> Following Goodman (1955, pp. 74-81).

“grue” does not, and that the time at which an object is first observed does not affect its color.

None of this probabilistic information can plausibly be construed as analytic in any relevant sense: it is not logically inconsistent, nor is it inconsistent with the axioms of probability theory,<sup>9</sup> nor does it somehow conflict with the meanings of words, to assign a higher prior probability to BIVH than to RWH, or to assign a higher probability to objects’ being uniform with respect to grueness than to their being uniform with respect to greenness.

## 2.2. *The Inferential Background Is Not Empirical*

So far, all of this is simply to say that an empirical inference requires some background information that is neither analytic nor directly observational. The empiricist could accept this, *provided* that this background can itself be justified on the basis of observations. The background information is not logically *entailed* by any observations. Perhaps, however, it can be derived from observations via non-deductive reasoning.

How might this work? It is unclear what empirical evidence I could possibly bring to bear, without begging the question, on the issue of how likely it is that I am a brain in a vat. But at least in the precognition case, I might cite the history of failed attempts to find psychic phenomena, the number of putative cases that have turned out to be hoaxes, the number of failed replication attempts, and so on. I might use this sort of inductive evidence to establish a very low prior probability for the particular ESP hypothesis entertained by Bem.

In the case of the shamrock inferences, perhaps I rely upon earlier experiences in which I used non-grue-like predicates to make inductive inferences, and the inferences turned out well. Or perhaps I have past experiences that inductively confirm that objects’ colors are generally not affected by time of first observation.

While these suggestions may seem reasonable enough, they fail to confront the underlying philosophical problem. By invoking further empirical inferences, they merely take the first step in a regress that the empiricist cannot complete. These further inferences will be subject to the same problem as the original inference: *every* empirical inference depends for its cogency on a suitable initial probability distribution. When *e* and *b* are logically independent (neither entails or contradicts the other), it is *always* true that a sufficiently low prior probability for *b* will leave *b* highly improbable in the light of *e*, and it is *always* true that

---

<sup>9</sup> Hereinafter, when I speak of the axioms of probability, I mean the Kolmogorov axioms: (i) the probability of any proposition is greater than or equal to zero, (ii) the probability of a tautology is 1, (iii) if A and B are mutually exclusive, then the probability that at least one of them is true equals the probability of A plus the probability of B, and (iv) the probability of A and B both being true equals the probability of A being true times the probability of B being true *given* that A is true. An assignment of probabilities is “coherent” or “probabilistically coherent” provided that it satisfies these four axioms.

there are coherent probability distributions on which  $e$  fails to raise the probability of  $b$ . So we would always need an independent justification for ruling out such probability distributions. Since we cannot have an infinite series of empirical reasons, the problem of the priors cannot in general be solved by empirical reasoning.

### 2.3. *The Straight Rule and Similar Proposals Fail*

Perhaps the dependence on a priori information can be avoided by adopting something like the straight rule: given  $n$  observations of  $A$ 's, out of which  $m$  turned out to be  $F$ , set the probability of the next observed  $A$  being  $F$  equal to  $m/n$ .<sup>10</sup> This rule putatively avoids the threat of a priori prior probabilities by specifying the correct credence in the inductive conclusion purely as a function of the observed evidence, without regard to how likely an  $A$  that is  $F$  may initially seem to the observer.

One problem with this rule is that it is inconsistent. Suppose I have just observed three shamrocks,  $s_1$ ,  $s_2$ , and  $s_3$ . The first and third were green, while the second was brown. What is the probability that the next two shamrocks I observe ( $s_4$  and  $s_5$ ) will both be green? There are at least two ways of answering this based on the straight rule:

1. Two thirds of observed shamrocks so far have been green, so the probability of  $s_4$  being green is  $2/3$ . If  $s_4$  does turn out to be green, the observed frequency of green shamrocks will then be up to  $3/4$ , at which point the probability of  $s_5$  being green will therefore be  $3/4$ . So the probability of both being green is  $(2/3)(3/4) = 1/2$ .
2. The frequency of observed *pairs* of shamrocks that are both green is  $1/3$ : out of the pairs  $\{s_1, s_2\}$ ,  $\{s_2, s_3\}$ , and  $\{s_1, s_3\}$ , only the last pair are both green. So the probability of the pair  $\{s_4, s_5\}$  both being green is  $1/3$ .

Perhaps there is a way of qualifying the straight rule to render it consistent. Or perhaps some other rule can be devised that, on its face, does not require us to employ any information beyond the sequence of observations. But here is a general problem with any such rule: either the results of the rule are consistent with the results of conditionalization starting from some coherent prior probability distributions (for short: “the rule is consistent with some priors”), or they are not. If the rule is not consistent with any priors, then the rule is incompatible with the axioms of probability theory and should be rejected for that reason.

If the rule is consistent with some priors, there will nevertheless be some *other* priors with which the rule is inconsistent. This must be true since, as

---

<sup>10</sup> A version of this rule is defended by Reichenbach (1938, pp. 340, 348-57) on quasi-pragmatic grounds; however, the defense fails to support the straight rule over many alternatives. For brief but effective criticisms of the straight rule, see Carnap 1980, pp. 85-6.

discussed above, different priors lead to different results in the light of the same evidence. Therefore, the empiricist will still stand in need of a justification for rejecting those priors that conflict with the empiricist's proposed rule. Once again, if the empiricist tries to provide empirical reasoning for that rejection, we will be embarked on a regress. The regress can only be stopped by appeal to a priori probabilities.<sup>11</sup>

#### 2.4. *Objections to Russell's Argument Do Not Apply to Mine*

The argument thus far is reminiscent of an old argument advanced by Bertrand Russell and later reprised by Laurence Bonjour.<sup>12</sup> Russell argued that inductive reasoning presupposes a *principle of induction* – a rule that would tell one, in effect, what conclusions may legitimately be inductively inferred from what evidence. The correctness of such a rule evidently could not be directly observed or deduced from direct observations. Nor could it be based on induction, on pain of circularity (of a sort that intuitively seems vicious). So the rule would have to be justified a priori.

Here I want to examine how an empiricist might respond to Russell, and then explain why that response does not apply to my own argument. An empiricist could reply to Russell that, even if there are rules that codify cogent induction, the subject who acquires inductive justification for a conclusion need not first know, or have justification for believing anything about, these rules. It suffices that the subject have dispositions to accept inferences that in fact follow the correct rules. To think otherwise – to hold that inferential justification requires first having justification to believe in the rules of inference – is to wrongly treat inference rules like *premises* of an inference.<sup>13</sup>

To elaborate on this response to Russell, compare two related points: First, in order to justifiedly classify some object as falling under some concept, one need not first have justification for accepting some analysis of that concept. Epistemologists, for example, have yet to determine the correct analysis of “know,” yet we all make justified knowledge-ascriptions all the time. In a similar manner, might we not be able to identify cases of cogent reasoning without knowing the rules describing such reasoning?

Second, in order to acquire justification for believing P via method M, it is not in general required that one first have justification for believing that M is reliable, or a source of justification, or in some other sense a good way of

---

<sup>11</sup> Thus, for example, even if Reichenbach could provide an epistemic reason for relying on the straight rule rather than any of the other coherent rules that allow learning from experience, he would not have vindicated an empiricist account of induction; he would rather have refuted empiricism by producing an instance of synthetic, a priori justification.

<sup>12</sup> Russell 1912, ch. 6; Bonjour 1998, ch. 7.

<sup>13</sup> Arguments of this kind appear in Psillos 1999, pp. 82-9; Huemer 2002, pp. 332-3; Devitt 2014. Psillos and Devitt both argue that it is permissible to use abduction to infer that abduction is reliable, because “rule circularity,” unlike “premise circularity,” is permissible in reasoning. On the problem of treating inference rules as premises, see also Carroll 1895.

acquiring beliefs. Any such general requirement would lead to an infinite regress, since one would have to acquire justification for believing that M is reliable using some other method, M', and one would then need justification for believing M' to be reliable. Now, if we reject the general requirement that one have justification for believing one's belief-forming method to be reliable, then there seems to be no principled reason for applying a version of the requirement to the special case where method M consists in a certain sort of inference.<sup>14</sup> For example, if justified perceptual belief does not depend upon one's having independent justification for believing that perception is reliable, then it is plausible to hold that justified *inductive* belief also does not depend on one's having independent justification for believing that induction is reliable (or cogent, etc.).

My purpose here is not to assess the merits of this reply. The point I want to make is that, whatever we might think of that reply as a reply to Russell or Bonjour, it could not be made in response to my argument. For I do not treat rules of inference like premises, nor do I claim that inferential justification requires prior knowledge about the rules of inference. I do not suppose that the rationality of some instance of reasoning turns on the subject's holding, or having justification for holding, any attitude to any propositions *about logical or epistemological relations*, or about the subject's own reasoning. Instead, I claim that the cogency of a piece of non-demonstrative reasoning depends upon a prior probability distribution, where the probabilities in question are probabilities of ordinary, non-logical, non-epistemological propositions – propositions such as “the 100<sup>th</sup> observed shamrock will be green” or “two thirds of all shamrocks are green.” The only relevant attitudes that the subject might need to hold, or to have justification for holding, would be credences in those ordinary propositions (N.B., not *beliefs about* credences). This is much closer to finding a priori *premises* on which empirical reasoning depends than to merely finding patterns of inference on which it depends.

### 2.5. *Skepticism Is an Unsatisfactory Response*

If the problem facing empiricism is that empirical reasoning seems to depend upon synthetic, a priori information, an easy response to the problem is to declare that there is no successful empirical reasoning, that non-demonstrative inferences never in fact expand our body of justified belief beyond direct observation. There is thus no need to embrace synthetic, a priori justification. This is the path of inductive skepticism.<sup>15</sup> Is this a reasonable way of preserving empiricism?

It is not. An epistemological theory should explain paradigm instances of

---

<sup>14</sup> Huemer 2002, pp. 338-9. Fumerton (1995, ch. 7) recognizes the problem but nonetheless maintains that inferential justification requires justification for endorsing the cogency of the inference.

<sup>15</sup> Hume [1758] 1975; Popper 1959.

knowledge and justified belief. Or perhaps more cautiously: a key desideratum of an epistemological theory is that it accommodate our pretheoretical convictions about what count as clear cases of knowledge, justified belief, good reasons, and the like.<sup>16</sup> Among these clear cases are many beliefs arrived at by empirical reasoning:

- If someone asks a chemist, “What is water composed of?”, the proper response is, “Water is composed of hydrogen and oxygen.” The proper response is not, “No one knows.”
- If someone asks a biologist, “Why should we believe the theory of evolution?”, it would be correct to reply that we have a great deal of evidence, drawn from such things as fossil records, physical similarities among organisms of different species, and so on. It would not be correct to reply, “We have no reason to believe the theory of evolution.”
- If someone asks an astronomer what cosmological theories can reasonably be held given our current evidence, it would be appropriate for the astronomer to mention the Big Bang Theory and brane cosmology. It would not be appropriate to include Aristotle’s theory of the planetary spheres, or the Native American theory that the Earth rests on the back of a giant turtle.

These are not difficult cases; the above judgments are entirely uncontroversial among scientists. It is in fact largely from the impressive accumulation of scientific knowledge of just that sort in the last few centuries that empiricism gained its popularity in the twentieth century. Whether rightly or wrongly, it was thought that scientific knowledge progressed as a result of its adoption of an empirical methodology and rejection of a priori intuitions, and empiricists sought to found all knowledge on the sources used by modern science.<sup>17</sup> It would therefore be irrational to jettison all of modern scientific knowledge in order to preserve the doctrine of empiricism. And it would be, if not quite contradictory, at least worrisome to treat empiricism in the manner traditionally reserved for a priori axioms – holding it fixed and giving up any and all particular judgments that are found to conflict with it. The empiricist, of all people, should stand prepared to revise his theory to accommodate the cases.

Note that there is no specially strong sense of “know”, “justified”, or “reason” in play here. The inductive skeptic is not merely one who denies

---

<sup>16</sup>This is (a more moderate version of) the view Roderick Chisholm (1982) defended under the name “particularism.” Chisholm held that epistemologists should begin from plausible epistemological judgments about particular cases and then seek to formulate general epistemological principles that account for those cases. As he notes, empiricism has generally rested on the opposite approach, “methodism,” which starts from general principles and uses them to arrive at judgments about cases. My view is more moderate than Chisholm’s, in that I claim only that accommodating the cases is an important desideratum for a theory.

<sup>17</sup> See again footnote 3 above.

absolute certainty to our scientific theories, or who denies the availability of non-circular metajustifications for our belief-forming methods, or who rejects some other idealized, philosophical sort of knowledge. (In this, the skeptic differs from the externalist, whose view is discussed in section 4 below.) Nor is there an option for the skeptic to hold that scientific theories are *probably* correct, *close* to correct, or even probably close to empirically adequate – for each of those claims would require empirical reasoning of just the sort the skeptic calls into question. Inductive skepticism, by definition, holds that there is no reason whatsoever for believing the conclusion of any empirical reasoning. On this view, there is no reason at all for preferring Copernican astronomy over “giant turtle” cosmology, no reason for favoring the theory of evolution over Creationism, and so on. This view is not consistent with the fundamental motives of empiricism, nor is it on its face a believable position.

### 3. Subjective Bayesianism Cannot Save Empiricism

#### 3.1. *Subjective Bayesianism and the Convergence of Opinion*

Subjective Bayesians hold that any coherent distribution of initial probabilities is rationally permissible.<sup>18</sup> Whatever one’s initial probabilities, one is rationally required, upon acquiring new evidence, to update one’s beliefs by conditionalization; that is, one should set one’s new credence in a given hypothesis equal to what one previously held as the probability of that hypothesis given the evidence:  $P_{\text{new}}(h) = P_{\text{old}}(h|e)$ .<sup>19</sup> On this view, empirical evidence can probabilistically justify a particular conclusion, not absolutely, but relative to a particular individual’s initial credences. These initial credences might be described as “a priori” in one sense, since they are not based on experience – but since there is no claim that they are in any sense objectively correct, or any more correct than any other coherent credences, we might consider this doctrine to be consistent with empiricism.

A leading objection to subjective Bayesianism is that the theory renders epistemic rationality and empirical reasoning too subjective.<sup>20</sup> On the subjective Bayesian view, whether and how much some evidence non-deductively supports some conclusion is relative to an observer and dependent on that observer’s psychology, as is the question of to what degree any proposition is overall justified. Call this “the Subjectivity Objection.”

In reply, subjective Bayesians offer the *convergence theorems*, which show that,

---

<sup>18</sup> de Finetti [1937] 1980; Savage 1954; Jeffrey 1992. Howson and Urbach (2006, p. 301), however, claim that the notion of rational priors is “incapable of being given any coherent or sustainable interpretation.”

<sup>19</sup> Jeffrey (1992, ch. 3) advocates a more general rule than conditionalization, which he calls “probability kinematics,” but the distinction is unimportant for our present purposes. Howson and Urbach (2006, pp. 80-85) hold that the rule of conditionalization applies only on condition that one’s conditional probabilities are unchanged when one learns new evidence.

<sup>20</sup> Jaynes 1968, p. 228; Gelman 2008.

under certain conditions, individuals with disparate initial probabilities for some hypothesis will tend to converge toward certainty in the true view as evidence accumulates.<sup>21</sup> For example, a pair of individuals with different initial opinions about the proportion of green marbles in a certain urn will tend to approach agreement (their subjective probability distributions will more closely resemble one another) as they progressively draw marbles from the urn and conditionalize on what they observe. More specifically, both individuals' probability distributions will come to resemble ever more closely the distribution that assigns probability 1 to the correct estimate of the proportion of green marbles.

### 3.2. *The Convergence Theorems Do Not Address the Problem of Subjectivity*

Even if we grant all the assumptions of the convergence theorems, the theorems do not touch the core of the Subjectivity Objection. To address the real subjectivity problem, what needs to be shown is something that these theorems do not even purport to show. Here is roughly what the theorems claim:

*The Convergence Principle:* Given a pair of initial subjective probability distributions, and given a desired degree of convergence in posterior probabilities in the light of evidence, it is possible to identify an amount of evidence that would induce the desired degree of convergence, starting from those initial distributions.

But that is perfectly compatible with the following:

*The Divergence Principle:* Given any set of evidence, and any desired degree of *divergence* in posterior probabilities in the light of that evidence, it is possible to find a pair of initial probability distributions that induce the desired degree of divergence, after conditionalizing on the given evidence.

Both the Divergence Principle and the Convergence Principle are true, across a wide range of relevant cases. But it is the Divergence Principle that matters; it is what underwrites the Subjectivity Objection, properly understood.

Why is it the Divergence Principle that matters? The goal, presumably, is to secure a kind of objectivity for what we pretheoretically regard as well-established scientific theories and other beliefs that are well-supported by empirical reasoning. To put it more clearly, the goal is to enable us to say, under certain conditions, that a certain conclusion is *objectively* well-supported by the evidence. Now there are two important points about this:

i) Objectivity is not secured merely by the possibility of some rational

---

<sup>21</sup> Savage 1954, pp. 46-50; Gaifman and Snir 1982; Hawthorne 2011. For a survey of convergence results, see Earman 1992, ch. 6.

observers agreeing on something. *If* we wish to understand objectivity in terms of agreement, we must require agreement by *all* rational observers.

- ii) At any given time, an observer or community of observers has a specific, limited body of evidence. Therefore, if we wish to say that some beliefs are objectively well-supported, we must be able to say that they are objectively supported by such a limited body of evidence. No actual belief is rendered justified by features of some merely potential future evidence.

Combining points (i) and (ii), we must be able to say: there are some conclusions and some (finite) sets of evidence such that, given that evidence, all rational observers must agree on that conclusion. That is what objectivity would mean. And that is precisely what the subjective Bayesian cannot say.

### 3.3. Why the Divergence Principle Holds

Why can't the subjective Bayesian say that? Let  $b$  be some hypothesis and let  $e$  be the entire body of evidence relevant to  $b$  that we have at some particular time. Assume that  $P(e|\sim b)$  is nonzero (otherwise, the evidence conclusively establishes  $b$  and the issue is trivial). Then the probability of  $b$  in the light of  $e$  can be determined from just two pieces of information: (i) the prior probability of  $b$ ,  $P(b)$ , and (ii) the likelihood ratio,  $P(e|b)/P(e|\sim b)$ . Letting " $p$ " stand for the prior probability of  $b$ , and " $L$ " stand for the likelihood ratio, we have the following simple theorem:<sup>22</sup>

$$P(b|e) = \frac{Lp}{Lp + 1 - p} \tag{Equation 2}$$

Figure 1 contains a graph of that equation, with six different possible values of  $L$ .<sup>23</sup> For each value of  $L$ , posterior probability (the probability of  $b$  in the light of  $e$ , shown on the vertical axis), is given as a function of the prior probability (shown on the horizontal axis). When  $L$  is greater than 1, the graph curves toward the upper left, indicating that the final probability is greater than the initial probability, and thus that the evidence supports the hypothesis. When  $L$  is less than 1, the graph curves downward to the right, indicating that the posterior probability is lower than the prior probability, and the evidence undermines the hypothesis.

Here is the important point: for any nonzero value of  $L$ , the function in Equation 2 maps the interval  $[0,1]$ , continuously and one-to-one, onto the interval  $[0,1]$ . For every possible value of the prior probability between 0 and 1,

---

<sup>22</sup> Proof: start with Equation 1 from section 2.1 above in the text, substitute  $(1 - P(b))$  for  $P(\sim b)$ , and divide both the numerator and the denominator on the right hand side by  $P(e|\sim b)$ .

<sup>23</sup> Courtesy of the Desmos equation grapher, <<https://www.desmos.com/calculator/yr6lojrkcl>>, accessed January 9, 2016.

there is exactly one corresponding value of the posterior probability, *and* for every possible value of the posterior probability between 0 and 1, there is exactly one corresponding value of the prior probability. In other words: you tell me what you want the posterior probability to be, and what you think the likelihood ratio is, and I can tell you a prior probability that delivers exactly that posterior probability. So: *if there are no constraints on the prior, then there are no constraints on the posterior.*

The preceding argument does not apply just to some special class of cases, or some particular model of inductive reasoning. Equation 2 is a simple, general theorem applicable to any hypothesis and any evidence, with the sole restriction that  $P(e|\sim b)$  is nonzero.

Bayesians who are otherwise good subjectivists sometimes suggest special constraints on which hypotheses one may assign a prior probability of zero to. For instance, perhaps one should not assign probability zero to any theory that has been seriously advanced in the scientific community.<sup>24</sup> These sorts of constraints make no difference to the above point. If you like, in the graph of Figure 1, delete just the points  $(x=0, y=0)$  and  $(x=1, y=1)$ . The graph then shows a set of one-to-one functions that map the open interval  $(0,1)$  onto the open interval  $(0,1)$ . It is then no longer possible to deliver a posterior probability of 0 by choosing a prior of 0, but it is still possible to deliver posterior probabilities arbitrarily close to 0. Given that every prior strictly between 0 and 1 is rationally permissible, no posterior strictly between 0 and 1 can be ruled out

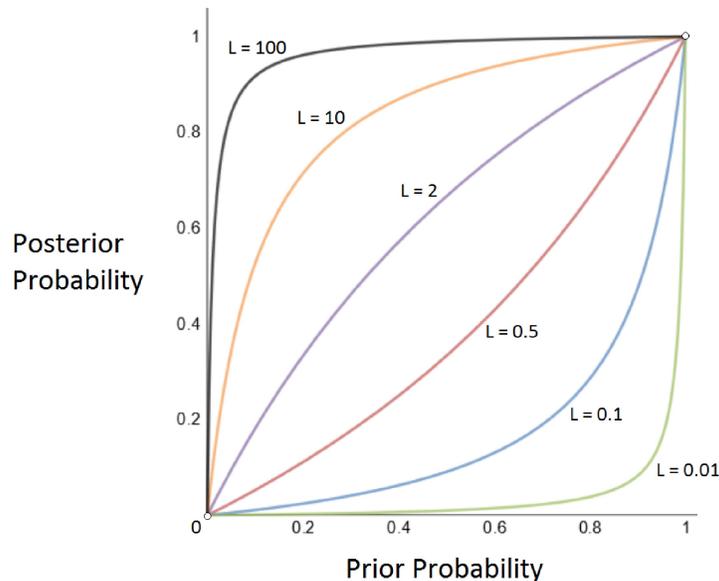


Figure 1

<sup>24</sup> Shimony 1970, pp. 97-103. The convergence theorems typically require all agents to initially agree on which propositions have prior probability zero.

for all rational agents.

### 3.4. *Evidential Relevance Must Be Subjective Too*

As Figure 1 shows, individuals with different prior probabilities for a hypothesis will disagree on the posterior probability of that hypothesis. But they might nevertheless agree on the *direction of change* of the probability. That is, they might agree, qualitatively, on whether the evidence probabilistically confirms, disconfirms, or is irrelevant to the hypothesis. In that case, the Bayesian might lay claim to some objective epistemic judgments, of the form “*e* is evidence for/against *b*,” albeit not of the form “*b* is justified.”

Bayesian reasoners will agree on evidential relevance in this sense, if and only if they agree on whether the likelihood ratio is greater than, equal to, or less than 1.<sup>25</sup> (They needn’t agree on the precise likelihoods, nor the precise value of the ratio.) So the question now is whether the subjective Bayesian is entitled to assume such agreement.

Here is one route to agreement. One might think that, in the case of many interesting scientific theories, the theory *entails* certain observational predictions. In that case, as all Bayesian reasoners will agree,  $P(e|b) = 1$ , where *e* reports such an observation. At the same time, the negation of the theory will fail to entail these same observational predictions. And one might think that, given that  $\sim b$  fails to entail *e*, all rational observers should agree that  $P(e|\sim b) < 1$ . Therefore, in these cases, all rational observers should agree that  $P(e|b)/P(e|\sim b) > 1$ .

There are two problems with this line of thought. First, it is unclear how the subjectivist can assume that all rational agents will agree on whether  $P(e|\sim b)$  is less than 1 or is instead equal to 1. It does not follow from the axioms of probability that when A fails to entail B,  $P(B|A)$  must be less than 1; indeed, it is easy to see that there must be many cases in which this fails.<sup>26</sup> So there is no pure probability-theory reason for rejecting the equation  $P(e|\sim b) = 1$ .

Second, and more importantly, interesting scientific theories virtually never entail observational predictions.<sup>27</sup> They merely render observational predictions plausible, sometimes extremely plausible, given our background probability distribution. Consider two examples:

---

<sup>25</sup> This is the Directional Agreement Condition assumed by Hawthorne’s (2011) convergence theorem.

<sup>26</sup> Suppose there is a continuous infinity of possible worlds in which A holds. Then it can’t be the case that each of these worlds has a nonzero probability. So take a world that has probability zero, and let B be the proposition that the actual world is *not* that world. Then A does not entail B, but the probability of B given A is 1.

<sup>27</sup> Compare Duhem [1914] 1954, pp. 185-7; Quine 1951, pp. 39-43, but note that whereas Quine and Duhem hold that a scientific theory conjoined with many other beliefs enables one to deduce observational predictions, my claim is that a scientific theory renders observational outcomes more or less probable, relative to our background credences. There is no need for any part of our belief system, including the whole of it, to *entail* any observational predictions.

- i) *Water is composed of hydrogen and oxygen*: How do we know this? One type of experiment involves burning hydrogen in the presence of oxygen. A certain mass of water is produced, while an equal mass of hydrogen and oxygen is consumed. Another experiment involves applying voltage to a sample of water. Bubbles of hydrogen and oxygen gas are produced at the negative and positive terminals, while the mass of water present decreases by the same amount as the total mass of hydrogen and oxygen produced.<sup>28</sup>

Neither of these phenomena is entailed by the theory that water is composed of hydrogen and oxygen. Even if water were composed of hydrogen and oxygen, it still could have been that we could find no way of separating water molecules into their constituent elements, nor of synthesizing the component elements into water.

Nor is the experimental evidence incompatible with alternative theories of the composition of water. The cited evidence is not incompatible, for example, with the hypothesis that water is a compound of sodium and uranium, since it *could* be that burning hydrogen causes a certain mass of hydrogen to transmute into sodium, while a certain mass of oxygen transmutes into uranium, and it could be that electrolysis induces a reverse transmutation.

- ii) *Human beings evolved by natural selection*: This is suggested by, for example, the fossil record. But the theory of evolution does not entail that there would be an available fossil record of human ancestors, or evolutionary ancestors of any other species. Nor do alternative theories of human origins logically contradict the fossil record. It *could* be, for example, that God created the Earth and all its species in the year 4004 B.C., and that He planted the fossils in the ground at that time.<sup>29</sup>

I think the alternative theories described above – the sodium-uranium theory and the Creationist theory – are objectively unjustified. If someone held those theories, and proposed the hypotheses suggested above to account for the empirical evidence, I would deem that person unreasonable.

But a subjective Bayesian cannot say this. The axioms of probability do not condemn anything about those theories. The axioms of probability only dictate the value of  $P(e|b)$  in the special case where  $b$  entails  $e$  or  $b$  entails  $\sim e$ : in the former case, one must assign  $P(e|b) = 1$ ; in the latter,  $P(e|b) = 0$ . In all other cases, any value between 0 and 1 inclusive is possible.

These are not unusual examples. The same point could be made for many other examples of scientific evidence: the observation of Foucault's pendulum as evidence for the rotation of the Earth; the matching of the shapes of the east coast of South America and the west coast of Africa as evidence for continental drift; the observation of the Chicxulub crater as evidence for the asteroid-impact

---

<sup>28</sup> Lavoisier 1783; Nicholson 1800.

<sup>29</sup> See the theory of Philip Gosse [1857] (2003).

theory of the extinction of the dinosaurs; and so on. Hardly any of modern science – perhaps none – can be described as objectively rationally supported on a subjective Bayesian account.

### 3.5. *Subjective Bayesianism Is a Form of Skepticism*

How bad is all of this? I think the subjective Bayesian's position is similar to that of the skeptic; indeed, subjective Bayesianism is a form of skepticism. Consider four paradigmatic epistemological judgments:

1. There are good reasons to think that water is composed of hydrogen and oxygen.
2. We are justified in believing that water is composed of hydrogen and oxygen.
3. Given current evidence, it would be irrational to think that water is composed of sodium and uranium.
4. We know that water is composed of hydrogen and oxygen.

All of those are uncontroversial in normal contexts. But on the subjective Bayesian view, a rational person apprised of all current scientific evidence could deny all of them. In that sense, none of them is objective. Furthermore, there is a case to be made that, given certain tenets of subjective Bayesianism, all of (1)-(4) must be false.

Subjective Bayesians would most likely wish to embrace (1)-(4). Perhaps they would adopt semantic theories according to which the terms of epistemic appraisal – “good reason,” “justified,” “irrational,” and “know” – have application conditions that are relative to the speaker's subjective probability distribution: if I happen to have a credence function on which [Water is composed of hydrogen and oxygen] is highly probable on current evidence, then I can correctly affirm (1)-(4). At the same time, of course, someone with a different credence function might correctly deny all of (1)-(4).

The question is how plausible such a view would be. Start with the idea of good reasons. Intuitively, there is a distinction to be drawn between what a person *is committed to* and what the person, all things considered, *has good reason to believe*. If I believe P, and P entails Q, then I am committed to Q. However, it does not follow that I have good reason for believing Q, since my initial belief in P might be purely arbitrary. If I believe P arbitrarily and without justification, it seems, then I do not, all things considered, have good reason to believe Q, nor am I *justified* in believing Q.

Subjective Bayesianism is more plausible as a theory of rational commitment than as a theory of justification or good reasons: if you have a certain initial distribution of credences, then it is plausible to hold that those initial credences commit you, on receiving new evidence, to the credences that result from conditionalizing on that evidence. But the claim that this process results in epistemic justification or good reasons is only plausible if the initial distribution

of credences was somehow justified. To hold that a purely arbitrary set of credences results in justification once we follow what that set of credences commits us to, is comparable to holding that a purely arbitrary premise justifies a conclusion, as long as one validly deduces the conclusion from the premise.

One might reject the latter analogy on the grounds that, in the Bayesian case, the subject would be conditionalizing on some genuine evidence, so the conclusion would not be reached *solely* on the basis of arbitrary credences. This is true. So a better analogy would be one in which, in a traditional, non-probabilistic setting, an epistemologist claims that if one validly deduces C from the conjunction of A and B, and A is justified but B is completely arbitrary, then C winds up justified.

Another reply on behalf of the subjective Bayesian: on the subjective Bayesian view, it is not that an unjustified starting point renders a conclusion justified. Rather, the initial credences are perfectly *justified*; it is just that many other (incompatible) starting points are equally justified.

But this is an extremely implausible view of justification. What could be conferring the justification in this case? The initial credences one starts with are explicitly said, in subjective Bayesian doctrine, to be no better than any other coherent set of credences. All that can be said for one's actual distribution is that it is *not ruled out* by the constraints of probability theory. For example, a good Bayesian reasoner might start with an extremely high confidence that there are eleven purple lions within one million miles of herself, and when asked to justify this, she might reply that no law of probability precludes that proposition's having a probability of 0.999999. This really cannot be described as an adequate source of justification.

A related problem is that, if I am a subjective Bayesian, I believe that I have no rational reason for preferring my own prior probability distribution over any other coherent distribution. But I also know that some of these other distributions would lead to radically different beliefs from my current belief system, given the same evidence. I must hold these other belief systems to be less well calibrated to reality, though an equally rational reaction to the identical evidence. It therefore seems that I must hold that I just got *lucky* by happening to have priors calibrated with reality.<sup>30</sup>

I conclude that subjective Bayesians have no plausible way of accommodating judgments (1), (2), or (3). Since knowledge requires justification, the subjectivists also cannot account for (4). If this is true, then there is an enormous range of paradigmatic epistemological judgments that the subjective Bayesian cannot plausibly accommodate.

#### 4. Externalism Cannot Save Empiricism

Epistemological externalists hold that factors external to a subject's mind can

---

<sup>30</sup> Related concerns are raised by Hanson and Cowen (n.d.).

affect whether and to what degree the subject is justified in a given belief. For example, whether a belief is justified might turn on whether it was formed by a process that *in fact* (whether the subject has reason to believe this or not) tends to produce true beliefs.<sup>31</sup>

Why think externalism might be true? One sort of argument concerns cases in which a subject suffers some cognitive malfunction but is unable to detect the malfunction. For example, a subject might have false memories (while having no reason to think the memories false) of having proved that P. In such a case, many intuitively judge that the subject lacks justification for P. Even arch-internalist Laurence BonJour conceded this point, thus apparently allowing an external factor – whether one’s memories are veridical – to affect one’s justification for belief.<sup>32</sup>

Though I do not find the case for externalism persuasive, adjudicating the internalist/externalist dispute is too large a task to undertake here.<sup>33</sup> Here, I shall merely argue that whether externalism is correct or not, there is no help for the empiricist. For the only form of externalism that explains the epistemological data that we want explained is a rationalist form of externalism.

The epistemological data consists of such facts (taken from section 3.5 above) as

2. We are justified in believing that water is composed of hydrogen and oxygen.
3. Given current evidence, it would be irrational to think that water is composed of sodium and uranium.

As we have seen, to accommodate such facts, we must constrain the range of rationally permissible prior probability distributions. How might an externalist justify such constraints? The externalist might argue that the “justified” or “rational” priors are simply those that reliably lead to correct conclusions in the light of the sort of empirical evidence that we humans tend to receive. Alternately, generalizing on the notion of reliability for outright beliefs, the externalist might propose that a justified/rational credence distribution is one that is *well-calibrated* (meaning, roughly, that the things to which one assigns a

---

<sup>31</sup> Goldman 1979. More sophisticated versions of reliabilism incorporate internal constraints on justification; however, such qualifications are not relevant to my argument in the text.

<sup>32</sup> BonJour 1998, pp. 124-9. BonJour also introduces the requirements that, to obtain a priori justification for a belief, one must have reflected sufficiently carefully, and one must have an adequate understanding of the concept of necessity. Hilary Kornblith (2004) has seized the opportunity to announce the death of internalism.

<sup>33</sup> I believe BonJour was wrong to concede the point about memory. For an extended defense of internalism in general, see [---]. On the role of memory in justification, see [---].

probability of  $x^0\%$  tend to be true about  $x^0\%$  of the time).<sup>34</sup> The externalist can then hold that our actual priors, which led us to accept modern science, are justified because they are in fact well-calibrated or because they in fact tend to lead us to true beliefs. As far as the externalist is concerned, it doesn't matter why they are well-calibrated or truth-conducive, nor does it matter if subjects lack independent reasons to believe that they are well-calibrated or truth-conducive.

Now leave aside the question of whether that is a good epistemological theory or not. The question is, is it an *empiricist* theory? Here the answer is simply no. An empiricist theory is one that accounts for all justification for substantive (non-analytic) propositions in terms of observation and inference from observations. An empiricist externalist is of course *possible* – for example, a reliabilist who holds that the only unconditionally reliable belief-forming processes are observation-based. But that simply is not the theory described above. The above-described theory holds that prior probabilities are justified in virtue of their calibration or truth-conduciveness; the idea of deriving these priors from observation plays no role at all in the account of their justification. It is a theory of *innate* justified beliefs.

On one account, suggested by Robert Nozick, the explanation for the reliability of our inductive reasoning is to be found in evolution: natural selection favored those who saw inductive reasoning as cogent, because in our evolutionary history, inductive reasoning usually led people to form correct beliefs, which promoted their survival and reproduction.<sup>35</sup> Adapting this view to a probabilistic framework, we can hypothesize that natural selection favored well-calibrated prior probability distributions. Again, this is a theory of something like innate knowledge (or rather, innately justified and calibrated credences), not a theory of knowledge derived from experience.

Of course, epistemological *internalists* would reject this as a possible account of how our prior credences are justified, since they would say that the reliability of our credences by itself does nothing to render them justified. But externalists could not say that. Nor could anyone who endorses the evolutionary account coherently deny justification to our prior credences. For the evolutionary account itself could be justified only if the theory of evolution in general is justified, which could be true only if the prior credences that led us to accept that theory (in the light of our current evidence) are justified. Nozick avoided self-defeat by embracing an externalist account of reasons, according to which the evolutionary account of the source of our priors is simultaneously an

---

<sup>34</sup> More precisely, one may define a well-calibrated probability measure as one with a low Brier score, calculated as the mean squared difference between (a) the probability assigned to a given proposition, and (b) the truth-value of the proposition, where a true proposition has a value of 1, and a false proposition has a value of 0 (but note that this way of measuring calibration combines the desiderata of reliability and informativeness). See Brier 1950.

<sup>35</sup> Nozick 1993, pp. 108-9.

account of why those priors are *justified*.<sup>36</sup> None of this gets us out of the dependence on non-observational sources of justification.

## 5. Rationalism

Having rejected empiricism, skepticism, and subjectivism, we are left with a rationalist view of probabilistic reasoning. On this view, there are nontrivial credences, or credence ranges, that are justified a priori, and these underwrite probabilistic inferences from observations. Different accounts, both internalist and externalist, can be given of how these credences are justified; the justification will count as a priori provided that it is not based on observation. My own preferred account, which I lack space to discuss in detail here, would appeal to necessary constraints on logical probabilities, constraints that, I claim, can be seen intuitively to be correct.<sup>37</sup>

Any rationalist view of probabilistic reasoning is open to well-known objections. Among the challenges is the sheer difficulty of plausibly answering, for nearly any proposition that we care about, the question of what its a priori probability would be.<sup>38</sup> For example, what is the initial probability, prior to obtaining any empirical information about the world at all, that water would be composed of hydrogen and oxygen?

I am not about to solve these problems here. This paper concerns the problems facing empiricism, not those facing rationalism. But I will conclude with a few remarks about how it might be rational to embrace rationalism in the face of such unresolved difficulties.

To begin with, it is worth noting that the rationalist does not need the strong thesis that, for any proposition whatever, there is a unique numerical value that is its a priori probability. The rationalist need only maintain that, for many interesting propositions, including, say, scientific theories, there is a *limited range* of initial credences one can reasonably hold.<sup>39</sup> For example, perhaps there is no particular number that is the unique rational a priori credence for the proposition that the Earth rests on the back of a giant turtle. But it also is not the case that the entire range of credences from 0 to 1 could rationally be held – it would not, for example, be rational to believe with 99% confidence that there is such a world turtle, prior to all relevant experience.

This view avoids the problem raised for subjective Bayesianism in section 3.3. Consider again the graph from section 3.3, but this time assume that the

---

<sup>36</sup> Nozick 1981, pp. 248-61, 264-8; 1993, pp. 108, 111.

<sup>37</sup> In [---], I defend a version of the Principle of Indifference, appealing to an explanatory priority relation to resolve some of the problems with interpreting the Principle. On the problems facing the Principle of Indifference, see van Fraassen 1989, p. 303; Howson and Urbach 2006, pp. 266-72.

<sup>38</sup> Ramsey 1931, pp. 160-7; Shimony 1970, p. 87.

<sup>39</sup> For present purposes, “limited” means having an upper bound strictly less than 1 and lower bound strictly greater than 0.

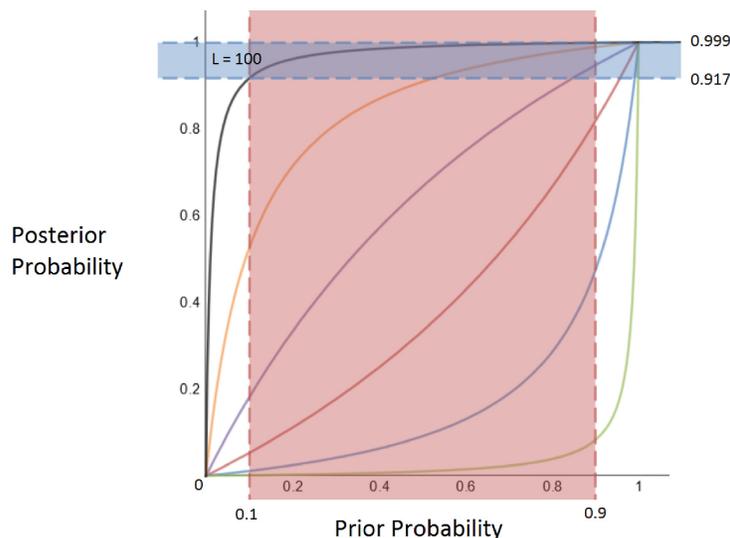


Figure 2

prior probability of the hypothesis, instead of being permitted to assume any value in  $[0,1]$ , is restricted to the interval  $[0.1, 0.9]$ . For illustration, focus on the case where  $L = 100$ . In Figure 2, I have added shaded rectangles corresponding to the restricted range of priors and the correspondingly restricted range of posterior probabilities. The allowed range of posterior probabilities is much narrower than the allowed range of prior probabilities: given that the prior is between 0.1 and 0.9, the posterior must be between 0.917 and 0.999.

This phenomenon is what the Bayesian convergence theorems are really about: in a wide range of circumstances, given a limited range of allowable priors, the range of allowable probabilities *narrows* when we conditionalize on evidence  $e$ . This does not help the subjective Bayesian, because the subjective Bayesian is not entitled to assume that the range of permissible priors is anything less than the full range from 0 to 1. But it does help a modestly objective Bayesian who holds that there are a priori constraints that limit to some degree the range of permissible prior probabilities.

This phenomenon also explains an otherwise puzzling feature of our probabilistic intuitions: we have more definite intuitions about the probabilities of certain hypotheses *in the light of our current evidence* than we do about their a priori probabilities. If someone asks me to estimate the a priori probability of the theory of evolution, I have very little to say. (I'm pretty sure it is not more than 90%. But I am not sure whether it is more than 1%.) But if someone asks me to estimate the current probability, based on everything we know, of the theory of evolution, I have a much better idea of the answer. (I am quite confident that one should say something strictly between 90% and 100%.)

The empiricist explanation: probability judgments depend upon empirical evidence for their justification. The questions about a priori probability are

harder to answer because there is no justification for any answer.

My explanation: rational probabilities become *more determinate* (that is, span a narrower range), the more relevant evidence we have. Typically, a wide range of subjective probabilities is rationally permissible to start with, but the range narrows as we gather more evidence. With a sufficient amount of evidence, the range of permissible credences approximates to a single value. As the range narrows, so our confidence that we can make a reasonable estimate increases.

All of this is by way of saying that the rationalist's commitments are considerably weaker than might at first glance appear. This does not resolve the difficulties raised for rationalism above, since I have not described how one may determine the acceptable ranges of prior probabilities.

This, however, is at best a weak reason for denying rationalism. If we have to choose between the propositions

1. There are rational and irrational (though coherent) ways of assigning prior probabilities, but no one has yet been able to systematically describe which ways are rational and which irrational.

and:

2. No one is justified in believing that water is made of hydrogen and oxygen, rather than sodium and uranium.

it is clear which of the two is less plausible. (2) is incredibly implausible. (1) is inconvenient but not particularly implausible. It is not especially surprising that we should have proved so far unable to adequately describe the rational constraints on assignment of prior probabilities. (Consider the range of other things that philosophers have had trouble adequately describing, from the meaning of "knowledge," to the principles of right conduct, to the justification for perceptual beliefs.) The difficulties critics have raised for objective Bayesians are serious and important difficulties. But they are not *strong evidence* for the *falsity* of the view, certainly not strong enough to rationally lead us to prefer a theory that engenders an extreme philosophical skepticism.

## References

- Ayer, Alfred Jules. 1952. *Language, Truth and Logic*. New York: Dover.
- Bealer, George. 1992. "The Incoherence of Empiricism", *Proceedings of the Aristotelian Society* 66, supplement: 99-138.
- Bem, Daryl. 2011. "Feeling the Future: Experimental Evidence for Anomalous Retroactive Influences on Cognition and Affect," *Journal of Personality and Social Psychology* 100: 407-25.
- BonJour, Laurence. 1998. *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- Brier, Glenn. 1950. "Verification of Forecasts Expressed in Terms of

- Probability,” *Monthly Weather Review* 78: 1-3.
- Carnap, Rudolf. 1932. “The Elimination of Metaphysics Through Logical Analysis of Language,” trans. Arthur Pap, pp. 60-81 in *Logical Positivism*, ed. A.J. Ayer. New York: Free Press.
- Carnap, Rudolf. 1980. “A Basic System of Inductive Logic, Part II,” pp. 7-155 in Richard Jeffrey, ed., *Studies in Inductive Logic and Probability*, vol. 2. Berkeley: University of California Press.
- Carroll, Lewis. 1895. “What the Tortoise Said to Achilles,” *Mind* 4: 278-80.
- Chisholm, Roderick. 1982. “The Problem of the Criterion,” pp. 61-75 in *The Foundations of Knowing*. Minneapolis, Minn.: University of Minnesota Press.
- de Finetti, Bruno. [1937] 1964. “Foresight: Its Logical Laws, Its Subjective Sources,” tr. Henry Kyburg, pp. 99-158 in Henry Kyburg and Howard Smokler, eds., *Studies in Subjective Probability*. New York: Wiley.
- Devitt, Michael. 2014. “There Is No A Priori,” pp. 185-94 in *Contemporary Debates in Epistemology*, 2<sup>nd</sup> edition, ed. Matthias Steup, John Turri, and Ernest Sosa. Malden, Mass.: Wiley-Blackwell.
- Duhem, Pierre. [1914] 1954. *The Aim and Structure of Physical Theory*, trans. from 2<sup>nd</sup> edition by P. W. Wiener. Princeton, N.J.: Princeton University Press.
- Earman, John. 1992. *Bayes or Bust?* Cambridge, Mass.: MIT Press.
- Fumerton, Richard. 1995. *Metaepistemology and Skepticism*. Lanham, Md.: Rowman & Littlefield.
- Gaifman, Haim, and Marc Snir. 1982. “Probabilities over Rich Languages,” *Journal of Symbolic Logic* 47: 495-548.
- Gelman, Andrew. 2008. “Objections to Bayesian Statistics,” *Bayesian Analysis* 3: 445-50.
- Goldman, Alvin. 1979. “What Is Justified Belief?”, pp. 1-23 in *Justification and Knowledge*, ed. George Pappas. Dordrecht: D. Reidel.
- Goodman, Nelson. 1955. *Fact, Fiction, and Forecast*. Cambridge, Mass.: Harvard University Press.
- Gosse, Philip. [1857] 2003. *Omphalos: An Attempt to Untie the Geological Knot*. London: Routledge.
- Hanson, Robin and Tyler Cowen. n.d. “Are Disagreements Honest?”, Mercatus Center working paper, <<http://mercatus.org/publication/are-disagreements-honest>>, accessed January 7, 2016.
- Hawthorne, James. 2011. “Confirmation Theory,” pp. 333-90 in *Philosophy of Statistics*, ed. Prasanta S. Bandyopadhyay and Malcolm R. Forster. Amsterdam: Elsevier.
- Howson, Colin and Peter Urbach. 2006. *Scientific Reasoning: The Bayesian Approach*, 3<sup>rd</sup> ed. Chicago, Ill.: Open Court.
- Huemer, Michael. 2002. “Fumerton’s Principle of Inferential Justification,” *Journal of Philosophical Research* 27: 329-40.
- Huemer, Michael. 2005. *Ethical Intuitionism*. New York: Palgrave Macmillan.
- Huemer, Michael. 2009. “Explanationist Aid for the Theory of Inductive Logic,” *British Journal for the Philosophy of Science* 60: 1-31.

- Huemer, Michael. 2016. "Inferential Appearances," pp. 144-60 in *Intellectual Assurance: Essays on Traditional Epistemic Internalism*, ed. Brett Coppenger and Michael Bergmann. Oxford: Oxford University Press.
- Hume, David. [1758] 1975. *Enquiry Concerning Human Understanding*. Reprinted in *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, edited by P. H. Nidditch. Oxford: Clarendon Press.
- Jaynes, Edwin T. 1968. "Prior Probabilities," *IEEE Transactions on Systems Science and Cybernetics* 4: 227-41.
- Jeffrey, Richard. 1983. *The Logic of Decision*, 2<sup>nd</sup> ed. Chicago: University of Chicago Press.
- Kornblith, Hilary. 2004. "Conditions on Cognitive Sanity and the Death of Internalism," pp. 77-88 in *The Externalist Challenge*, ed. Richard Schantz. Berlin: Walter de Gruyter.
- Lavoisier, Antoine. 1783. "Report of a Memoir Read by M. Lavoisier at the Public Session of the Royal Academy of Sciences of November 12, on the Nature of Water and on Experiments which Appear to Prove that this Substance Is Not Strictly Speaking an Element but that it Is Susceptible of Decomposition and Recomposition," *Observations sur la Physique* 23: 452-5. English translation by Carmen Giunta available at <<https://web.lemoyne.edu/giunta/laveau.html>>, accessed January 9, 2016.
- Nicholson, William. 1800. "Account of the New Electrical Apparatus of Sig. Alex Volta, and Experiments Performed with the Same," *Journal of Natural Philosophy, Chemistry and the Arts* 4: 179-87. Available at <<https://books.google.com/books?id=HdkPAAAAQAAJ&pg=PA179>>, accessed January 9, 2016.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, Mass.: Harvard University Press.
- Nozick, Robert. 1993. *The Nature of Rationality*. Princeton, N.J.: Princeton University Press.
- Popper, Karl R. 1959. *The Logic of Scientific Discovery*. New York: Basic Books.
- Psillos, Stathis. 1999. *Scientific Realism: How Science Tracks Truth*. London: Routledge.
- Quine, W.V. 1951. "Two Dogmas of Empiricism," *Philosophical Review* 60: 20-43
- Ramsey, Frank. 1931. "Truth and Probability," pp. 156-98 in *The Foundations of Mathematics and other Logical Essays*, edited by R. B. Braithwaite. London: Kegan, Paul, Trench, Trubner & Co.
- Reichenbach, Hans. 1938. *Experience and Prediction: An Analysis of the Foundations and the Structure of Knowledge*. Chicago, Ill.: University of Chicago Press.
- Rouder, Jeffrey N. and Richard D. Morey. 2011. "A Bayes Factor Meta-analysis of Bem's ESP Claim," *Psychonomic Bulletin and Review* 18: 682-9.
- Russell, Bertrand. 1912. *The Problems of Philosophy*. New York: Oxford University Press.
- Savage, Leonard. 1954. *The Foundations of Statistics*. New York: Wiley.
- Schlick, Moritz. 1974. *General Theory of Knowledge*, tr. Albert Blumberg. New

York: Springer-Verlag.

Shimony, Abner. 1970. "Scientific Inference," pp. 79-172 in Robert G. Colodny, ed., *The Nature and Function of Scientific Theories: Essays in Contemporary Science and Philosophy*. Pittsburgh, Penn.: University of Pittsburgh Press.

van Fraassen, B. 1989. *Laws and Symmetry*. Oxford: Clarendon Press.