

THE FALSE HOPES OF TRADITIONAL EPISTEMOLOGY

Bas C. van Fraassen

(published version = *Philosophy and Phenomenological Research* 60 (2000), 253-280.)

After Hume, attempts to forge an empiricist epistemology have taken three forms, which I shall call the First, Middle, and Third Way. The First still attempts an a priori demonstration that our cognitive methods satisfy some (weak) criterion of adequacy. The Middle Way is pursued under the banners of naturalism and scientific realism, and aims at the same conclusion on non-apriori grounds. After arguing that both fail, I shall describe the general characteristics of the Third Way, an alternative epistemology suitable for empiricism.

I. The Death of Modern Epistemology

Philosophers always rewrite their own history. Hans Reichenbach did so in 1947, at a high point in his career (Reichenbach 1948). Any such retrospective is of course itself a philosophical tactic, thinly disguised: a history of past errors presented as leading us inexorably to the inescapably right philosophical position. Well, may it do so! After a brief look at Reichenbach's historical view I'll offer a revision and then complete the picture as I see it.

1.1 Death by dilemma: The disaster of "classical" empiricism

Reichenbach begins with the Rationalist-Empiricist strife of the 17th and 18th centuries.[1] With mathematics as paradigm of knowledge, the Rationalist imposed a simple requirement on philosophical accounts of science: show that it brings us a demonstrably true theory of the world. The question "How is knowledge possible?" is understood as "How is it possible to have indubitably true empirical knowledge?" The answer could only be that science must rest on indubitably true *a priori* principles.

Empiricists rejected the presupposition of the question as well as the answer. *A priori* reason is empty; it can only establish connections between ideas and these entail nothing about contingent matters of fact. That conclusion denies the Rationalist answer; but how much of the question was retained?

Reichenbach's overview covers familiar ground: the quick progression from Bacon's optimism, via Newton's and Locke's unshakable faith that the scientific method is known and leads to truth, to Hume's skeptical disaster. What all agree upon as scientific knowledge, in the 18th century, goes far beyond the evidence of observation. Since logic and mathematics, being forms of *a priori* reasoning, cannot lead us from one to the other, science must follow *ampliative rules* to create this knowledge. Those are the rules, sung by Bacon and invoked by Newton, though never satisfactorily formulated, that jointly constitute *Induction*, broadly conceived.

An ampliative rule is a rule of inference, or rule for changing or updating our opinion, that goes beyond purely logical assimilation of the evidence. Hence the contention that these rules lead science to the truth cannot be demonstrated *a priori*. Nor can that contention be derived by Induction (which consists simply of the ampliative rules which we accept) on pain of circularity.

The conclusion is inescapable. Reliability of Induction cannot be part of our scientific knowledge. The Empiricist theory of science has run itself into the ground.

As Reichenbach sees it, Empiricists had not rejected enough of the Rationalist's question. Hence they found themselves "pushed into the hopeless position of proving that empirical knowledge was as

good as mathematical knowledge".[2] With David Hume, the hopelessness became blatantly obvious, and skepticism seemed to be the only alternative to rationalism. Thus, as Reichenbach depicted it, classical Empiricism shared Rationalism's stringent *criterion of adequacy*: that an epistemology must show how absolutely reliable knowledge is possible. The way out will be to reject this criterion.

1.2 Life after Disaster: Three Ways

Empiricists after Hume, including Reichenbach himself, did by and large accept *weaker* criteria for epistemology. William Whewell said proudly that no truly consilient theory had ever turned out to be false -- consilience was the hallmark of theories formed by Induction and Newtonian science was his prime example of such a theory. Peirce claims, only a little more modestly, that "Quantitative induction always makes a gradual approach to the truth, though not a uniform approach"(Peirce, volume 2, p.770). Reichenbach himself makes the claim still a little more modest:

It can be shown that if it is possible at all to make predictions, the inductive method is an instrument to find them. (op.cit. p.146)

Doesn't it seem that there is still a fairly strong criterion of adequacy attempted -- and claimed -- to be met? Here is a tentative formulation: I'll make it as weak as I can.

An epistemology must imply that, and show how, epistemic security is humanly attainable by the methods of the sciences, under favorable conditions, and that it is in fact attained to some reasonable degree. *Security* means here, possibly knowledge, perhaps certainty, but at least reliable and accurate beliefs and opinions.

As stated the criterion is still somewhat vague; it is not clear how much leeway we have gained. But the philosophical tactic is clear enough: we will settle for a weaker conclusion, which can indeed be demonstrated, and which is arguably sufficient to provide epistemology and scientific methodology with a solid foundation. I shall call the conviction that this is possible, and its pursuit, the *First Way* for empiricists after Hume's skepticism.

A closer look at Reichenbach shows that he actually vacillated between *three* main options that emerged in the 19th and 20th centuries. The First Way was certainly presaged in the 19th century, as the idea of what Larry Laudan called the *Self-Corrective Thesis*. I will elaborate on this, and on Reichenbach's version, below. By the end of that century, however, we see the emergence of two further Ways which reject even the above weaker criterion. One, which I'll call the *Middle Way*, we can associate with the philosophical turn to which Dilthey gave the name *Naturalism*. It has taken various forms. The form which I'll call *Moorean Scientific Realism* still shares with both the initial predicament and the First Way an important presupposition: that science creates scientific knowledge by following ampliative rules, and that we can know these to be reliable. There is a difference. The First Way contends that we can *show* that ampliative method to be a reliable means to security in some important sense. The Middle Way only insists that we *know* that it is so, as part of our scientific knowledge (which is of course attained by that very means).

But besides these two Ways there is a Third, more extreme, with which Reichenbach (and Peirce) also flirted a little. In the remainder of this Part and in the second Part I will argue that the First and Middle Ways do no better than the ideas so ruefully contemplated by Hume. Then, I hope, the Third Way will fall on more sympathetic ears.

1.3 The Self-Corrective Thesis

If science follows ampliative rules when it arrives at conclusions from the evidence, and it is not possible to demonstrate that these rules lead to the truth, what can we do?

A new conviction increasingly operative in 19th century theory of science is that this rule governed activity is *self-corrective*, and must, as a result of this continual self-correction, gradually lead to the truth, at least in the long run. Larry Laudan (1981) called this the *Self-Corrective Thesis*. That this conviction could have felt satisfactory can only be understood, I think, in the context of general optimism fed by the successes of natural science. (In the humanities, exactly the same picture of inquiry led scholars instead to confront the hermeneutic circle and the possibility of non-convergence -- a rather more pessimistic reflection.)

Larry Laudan has told us the story of how the idea of Self-Correction fared in the hands of Peirce. (There are also quite different accounts of the matter; see Levi (1980)). Peirce expressed the conviction that through systematic self-correction scientific inquiry would lead to the truth. He also saw that one could not satisfactorily leave this as a mere article of faith. The inductive method as he saw it, consists in systematic testing by well-designed experiments (cf. Peirce, *op. cit.* vol 2, p. 775). But this method comprises three ampliative rules: *crude* induction, which is mere universal generalization on positive instances; *qualitative* induction, more or less what was later called the hypothetico-deductive method; and *quantitative* induction, which is statistical inference based on sampling procedures. Are these really self-corrective methods?

Qualitative induction, the third ampliative practice, includes hypothesis testing as well as 'abduction' (akin to what we now call 'inference to the best explanation'). Peirce held this to be of vastly greater importance in science than the other two. But he could not show to be self-correcting in the requisite sense. Just like crude induction it may be self-correcting in the sense that new evidence will force us to give up earlier conclusions. This may then force us to look for *other* hypotheses and *new* explanations. Peirce could see no way to demonstrate that the ones we then come up with will be increasingly better in some concrete sense (see especially Laudan, *op. cit.*, p.239).[3] Nor could he show that if this process is carried out under varying historical circumstances, the results will converge!

1.4 Reichenbach's version of the Self-Corrective Thesis

When Reichenbach takes up the banner he mounts a heroic effort, based on two rather outrageous ideas: that probability is to be equated with long run frequency, and that all the ampliative methods used in the sciences reduce to probability assessment through numerical induction. *Au fond*, the only ampliative rule ever genuinely used or needed is the 'straight rule' which assigns probabilities on the basis of frequencies observed so far. Therefore, if *this* rule can be vindicated or justified, the entire scientific methodology is saved.

Even if we reject the basic ideas, his effort is instructive. In the first place, it brings out clear empirical conditions under which no rule, however self-corrective, can succeed. Secondly, it makes precise Peirce's intuitive notion of how self-correction 'must' lead to truth in the long run. But thirdly, by running into a clear, indeed elementary, problem which is unsolvable within its own framework, it brings to light a challenge pertaining to any approach within what I call here the First Way.

To illustrate Reichenbach's reasoning, imagine a weather forecaster who must each evening

announce a probability for rain the next day. He has his own rule for doing so, which may draw on any facts ascertainable by evening. (He does not have a crystal ball that shows the future!) If his is a good rule, under the conditions in which he finds himself, his probabilities will come increasingly closer to matching the actual frequency of rain. (There are various ways to make this precise; I'll return to that.)

Reichenbach's first point is that the conditions may actually be such as to defeat *any* rule this forecaster could have. For the relative frequency of rain in this sequence of days may not converge to a limit, it may swing ever more wildly up and down. The possibility of such divergence in nature suffices to refute, not only Bacon's and Newton's convictions, but also Peirce's more modest optimism. No rule, not even any self-correcting procedure, *must* lead to the truth in the long run.

Some rules will latch on to the actual regularities more quickly than others; this too is a relation between the rule and the conditions under which it is being applied. So, should we just hope to be lucky and get a forecaster whose rule happens to match our local weather? Has science just been lucky in its choice of ampliative methods, with respect to the bit of cosmic history in which we find ourselves? Here comes Reichenbach's version of the Self-Corrective Thesis: no rule is necessarily successful, but there is a rule which will succeed if any rule could succeed. That rule is the very simplest one, 'crude numerical induction', the self-corrective method par excellence. Working by that rule the forecaster announces as his probability for rain the next day precisely the proportion of rainy days among those observed so far. Now, if those proportions tend to a limit, so will his forecasts; and his forecasts will come increasingly close to matching that limit -- obviously!

In fact, even Reichenbach's proof that the straight rule must lead to truth in the long run *if any rule would do so* does not really prove what it seems to prove. At first sight, his argument is so short and simple that there can be nothing wrong with it. But we would all laugh at the forecaster who does this, for he will soon be announcing probabilities that hardly differ from day to day, winter and summer, spring and fall.

Of course Reichenbach noticed this, and pointed out that the forecaster can answer questions of conditional probability. Asked for the probability of rain in July s/he will reply with the proportion of rainy days among July days so far. But this is not enough. On the third of July we will ask a forecast for the next day. Should the announcement now be that conditional July probability? But the next day is Independence Day -- should he not announce his conditional Independence Day probability? Here we have come up against the 'problem of the reference class', as it came to be known.

Undoubtedly the forecaster can come up with some amendment to his rule which will tell him what to do. We should now try to make precise how we shall assess his performance -- and then ask again: is there a rule which will do as well as any rule will do, under any conditions? That is, we must ask the same question again, and the answer will now not be so trivial.

What happened in this little contretemps is that we challenged the forecaster by pointing to a specifiable *subsequence* of the total sequence. We want him to be 'calibrated' (that is, to have his forecasts tending toward a match with the frequencies) on that subsequence as well. We should be allowed to specify any salient subsequence in such a challenge. There is general way to do this, without knowledge of the calendar or the national holidays, etc. (and without knowledge of what factors he counts as significant for his forecasts). Let's specify as important subsequences the classes:

$S(r) = \{\text{days for which the forecast probability of rain was } = r\}$
for each number r . The desirable feature is clearly that the frequency of rain in $S(r)$, if it be non-empty (or if you like, infinite) should tend to limit r . We can be more easygoing and use small intervals rather than precise numbers, or some other way to reflect good but less than perfect accuracy. Some of our

forecasters may have good rules, as it happens, by luck or inspiration: but is there a rule that can be shown to do (in the long run) as well as any of them can do? To answer this, we need at once more precision and more generality.

1.5 A more precise formulation -- and refutation

This level of precision was achieved in statistics. [4] First of all, we must target a wide range of rules -- more generally, *forecasting systems* -- which assign probabilities solely on the basis of frequencies observed up to the point of forecasting. Frequencies are the sole basis: this is the hallmark of numerical induction. But we must allow much more. Forecast systems may draw on background knowledge or information assumed but available prior to the forecast. It may also allow for self-correction: when observation shows that a simple rule followed so far has done badly, there may be a clause which replaces that simple rule by another one. The replacement clause may allow (putative) learning from experience: it may have a definite recipe for change, based on an evaluation of which new rule would have done better in the past -- a touch of abduction here!

This is liberal, but there is also a constraint. The rules must be humanly usable. That is not easy to make precise, so the constraint adopted is a precise but weaker one: *computability*. There is little point in deriving results for rules which could be of use only to angels or Maxwellian demons.

Secondly, we must tighten up the evaluation of accuracy. What corresponds to truth, and to gradual approach to truth? Restricting ourselves still to numerical induction, the first requirement is of course that the forecast probabilities should converge to the actual long run frequencies. But as in the challenge to Reichenbach's straight rule, we should imagine an adversary who challenges rules which have been performing well over all. The challenge consists in pointing to specifiable subsequences and raising the suspicion that the rule is not performing very well *there*.

Again, however, the challenge must be humanly formulable: we must disallow adversaries aided by crystal balls, angels or demons. The minimal requirement along this line is that the challenge must point to a *computable* subsequence. The computation may be based at each point on what has happened so far. For example, I may suspect that a forecast rule for rain will perform well over all, but badly if we look just at those days which follow immediately on five rainy days. The factors drawn on may include any factors on which the forecast rule itself can draw. If the rule stands up to all these challenges in the long run, it is called *computably calibrated*. It is our most stringent criterion of correctness that still respects some human limitation.

This theory of forecasting systems is numeral induction extended, liberalized, mechanized, and motorized. What can it do? Well, as Reichenbach had noted, nature could be giving us a sequence in which the requisite relative frequencies do not converge to any limit. Then there exists no correct forecasting system at all.

However, if all the requisite limiting relative frequencies exist, then there is a correct forecasting system and it is essentially unique.[5] This can be demonstrated for computable rules -- again, angels are left out of the picture. Does this give us hope? The hope was formulated as a certainty by Reichenbach: with numerical induction properly tuned we have a method that will lead us to the truth *if any rule will*:

It can be shown that if it is possible at all to make predictions, the inductive inference is an instrument to find them; and the inference is justified because its applicability represents a

necessary condition of success.

(*op. cit.*, p.146)

Is this true, as we have now come to construe it? No, it is not.

Given the convergence of computably calibrated forecast systems -- if any exist -- can there be a 'universal' forecast system which will be correct in the long run -- if any system is? The answer is *No*. No matter how well we design our procedure for self-correction in response to the evidence, Nature may confute it.[6] The proof uses the form of diagonal argument now familiar to us from all areas where computability plays a role.[7]

Philosophers and scientists who cherished the ideal of Induction had already progressively lowered their sights. All that Reichenbach still claimed was that the lowliest and least fruitful of all inductive methods would, through its self-corrective procedure, demonstrably lead to the truth in the long run, if any rule could. (Nothing comparable was claimed to be demonstrable for the more innovative methods that could enrich science theoretically.) But even that final modest claim turned out to be untenable.[8]

II. The Middle Way, Realistically Pursued

All Hume proved, Reichenbach said dismissively, was "that the rationalist aim of knowledge is unattainable by empiricist methods" (*op. cit.*, p.142). But Reichenbach too tried to demonstrate something a priori: *that we have the means to attain the truth, with a guarantee of reaching it if it can be reached at all*, although without the ability to tell with certainty when we have arrived. He failed; but in fact history had moved on already, so to speak, before that weaker conviction of demonstrable security had been dashed.[9]

2.1 The Naturalistic turn

In the type of philosophical positions which Dilthey denoted as *Naturalism*, the very idea of an *a priori* foundation for science (methodological or epistemological no less than metaphysical) is explicitly rejected. Science itself is what we know, and philosophers as much as anyone else should proceed on the basis of what we know. It is mistaken as well as self-defeating for philosophers to discuss science from a vantage point independent of or prior to their actual historical situation, replete with the science of their day.

This is anti-Foundationalism become proudly self-conscious, and in that aspect I applaud it. Naturalism is characteristically perspectival: the conclusions we shall reach follow for us, given our historically conditioned epistemic position, and need not be cogent for all logically conceivable inquirers, thinkers, or cultures. Nevertheless, we are rationally within our rights to claim those conclusions as knowledge. Indeed, to say otherwise is to make an impossible demand for foundations or justifications which are in fact demonstrably beyond our reach.

But Naturalism is not the only option when that impossible demand for foundations is left behind. Secondly, Naturalism itself is not monolithic, it admits of degrees and varieties. One variety is especially relevant to our inquiry into ways out of the disaster empiricism suffered in epistemology. I shall call it "Moorean Scientific Realism" -- Moorean because it insists on not leaving our common sense and common convictions behind.

Like the First Way, this view retains the assumption that science creates knowledge by following ampliative rules, rightly called Induction, Abduction, and the like -- and that we must show these rules to be reliable. The difference is only that our epistemology may now stand on our scientific knowledge, rather than just on thin air.

2.2 Moorean Scientific Realism

The view I wish to discuss can be illustrated with a considerable literature on the naturalizing of epistemology and of philosophy of science. [10] But let me try to present it in my own way. We begin with a naturalistic reassessment of what philosophers are to do. We are to clarify and illuminate the possibility of scientific knowledge, the efficacy of scientific methods. This task is to be broadly conceived; for science and its methods are, in their principles, not foreign to our daily commerce. But we are to do this without leaving ourselves behind, so to speak, without attempting to step out of our historical situation and view it from nowhere. We ourselves have knowledge, painfully but gloriously accumulated through science. From within this vantage point, drawing on the knowledge we have, we can entertain and answer questions about human existence. These questions include how humans in general gain their reliable opinion, from veridical experience, by reliable methods. Any such question is to be answered on the basis of what science tells us about nature, to which we, after all belong. This is Moorean Scientific Realism.

As in the First Way, the criterion of adequacy is still to show that, and how, epistemic security is humanly attainable, and in part largely attained, through science. But now the task is to show that on the basis of our shared scientific knowledge.

What will be the claim made by the Moorean Scientific Realist? I formulate it as follows:

"Here is what we can confidently say on the basis of scientific findings. First of all, science does proceed by a certain method, which we have been able to specify, and which involves ampliative rules, namely the following: Secondly, on the basis of what we now know, it is no surprise that science has attained great success by means of this method. For accepted scientific knowledge implies that this method is a very good means to its end, under conditions which we have found to be the actual conditions in our universe. Thus we can explain our success so far, and can confidently predict, with high probability, that the method will prove reliable in the future."

I had to leave a blank; when it comes to specifying and formulating those ampliative rules we are all in a bit of a quandary as yet. In discussion we'll have to rely on the traditional notions about what they are like.

To sum up then: according to this view, the epistemologist's task is to formulate an account of our epistemic life which relies on the deliverances of science and its accepted theories to show how epistemic security is humanly attainable. Indeed, given its reliance on the scientific tradition, it had better, for coherence's sake, imply that this security is actual. Thus a certain view of the human condition is presupposed, in the conviction that science supports this presupposition. Well, let's look into this! Despite its prima facie air of question-begging circularity, the stated criterion is actually not trivial -- as we shall see.

2.3 Science Does Not Support the Presupposition

I will now argue that Moorean Scientific Realism is based on a mistake -- not based on science but on a confused medley of science and preconceived ideas. I see it as a commendable attempt to eschew Foundationalism, but not as one that succeeds. My argument has three parts. First, science tells us that we have much *relevant ignorance*. Second, science does not provide us with a measure of that ignorance, which we are left to supply for ourselves. Third, the special sciences, which investigate the conditions under which we actually live, operate on an *approximation principle*. This is in effect an assumption about how favorable those conditions are, and it is an assumption they cannot support without circularity.

2.3.1 Inductive and Abductive Methods Are Unreliable

There is, as far as I can see, exactly one shining ubiquitous example of induction and abduction: racism, ethnic prejudice, and stereotyping. "Those people don't want to work, they only want handouts" - challenge this sentiment and the speaker will say that he has met a lot of them and they are all like that. And he'll go on from this induction to his next theme with a rhetorical question "Some races or cultures are very disadvantaged and exploited by others -- but isn't there some reason why they were like that in the first place? Why they didn't develop a technology or civilization!" The question invites inference to the best explanation, of course.

The role of science here is not to endorse but to criticize and correct this reasoning. The critique has to include the admission that such reasoning is reliable under special circumstances. No matter how badly you think of it, you must probably admit that racial and ethnic stereotyping, with its fast fight or flight responses, saved lives in wars and riots, at the cost of sometimes doing individuals an injustice. In a nutshell, this is the scientific verdict: humanly usable ampliative methods for prediction are entirely unreliable in general, while reliable under historically important special circumstances. Prevalence of such methods is to be explained by the brute accident of their survival value under *special* circumstances in which we humans have lived and to which they happen to be suited.

In the scientific world picture it cannot be otherwise, for the humanly available indicators are at once gross and meager. Let's look at some examples. Weather forecasting is still based -- and must be based -- on such indicators, but gets us only a few days worth of reliable information, even with satellite photos. The weather has several much further reaching analogues.

For example, there is also 'cosmic weather'. A letter to the New York Times [11] points out that about 3000 meteors of significant size hit the earth each day. Assuming current recent figures of air traffic, the letter calculates that there is a 10% probability that one of these hits an airplane within a 30 year period. Believe me, the airline pilots receive no useful meteorological predictions on that subject! Yet cosmic weather matters: the dinosaurs were rulers of the Earth, but succumbed to it.

Similarly we can speak of microphysical weather. Predictability is limited to laboratory conditions. The metals, salts, and carbon compounds are normally not pure: microphysical storms rage all around and inside us. From time to time the microphysical weather deteriorates, and we die of unknown, hardly classifiable causes.

Let me return for a moment to ordinary weather forecasting. This has greatly improved during the last half century. We are rather lucky here on Earth in that the forecast rule "Tomorrow same as today" really does not do all that badly. That is the accepted performance baseline. It's what

meteorologists mean by "just guessing", though logically it is not at all random, it is a true ampliative rule. Today the papers publish 5 day forecasts, because no model we have, using the information we can obtain, does better than 'just guessing' on the 6th day.[12]

The problem is not that we do not understand the mechanics of weather. We think we do. But theoretical studies show that even under the perfect model assumption (assumption that the equations are exact) we get improvement over guessing for at most two weeks. And the improvement over guessing, even for tomorrow (let alone the penultimate forecast day!) is not phenomenal.

We are lucky that we happen to live in a rather stable environment, in which microphysical, cosmic, and earthly weather are all rather benign. On the assumption that our science is true, this luck, rather than the virtues of our inductive habits, is what saves us.

Looking at induction and similar 'methods', we tend to focus on the successful examples, and on the rationale people so often give for what they believe. Both are misleading. If we use induction (generalization from known examples, extrapolation from observed frequencies) it sometimes works and sometimes does not. Can induction tell us when this sort of extrapolation will succeed and when it won't?

This is the place where science has something to tell us: if science is true, success will depend on facts of microstructure and cosmic structure which cannot be among the input for human induction. So the answer is No: induction cannot tell us which applications of induction will succeed.[13]

I can easily imagine two objections. First of all, do we not enjoy a great deal of regularity and predictability? Think of the seasons and of cooking, two cases in which our record of successful prediction spans millennia. Secondly, haven't we learned to use our gross and meager indicators to gauge underlying conditions, on which to base further prediction? Temperature, tongue color, and pain serve to diagnose illnesses with predictable courses and response to medicine.

I grant both points. However, I see the second as but an instance of the first, and see both as fortunate regularities that we have fortunately latched onto. My focus on salient unpredictability was a device to highlight the part of fortune which is equally important to long-range and short-term prediction. But to support this view, I will clearly need to examine the question on a more theoretical level.

2.3.2 Science Does Not Tell Us the Boundary Conditions, Nor Give Us a Probability Measure over Them

To focus the discussion, despite its abstraction, let's assume for now that Newton's System of the World is the true science. What I shall argue is that if this is taken as the basis, in principle, of all science then we are not justified to expect its future success. Of course we may be rational to expect that nevertheless, but we cannot support the expectation on the basis of science (alone). The argument will apply *mutatis mutandis* to later candidates for basic science.

As Eugene Wigner observed, the crucial turn to modern science, completed by Newton, was the search for laws of succession: principles that govern how given conditions evolve into other conditions. Newton's important contribution was not that there are exactly six planets, or that planets follow elliptical orbits; but rather that *if* initial and boundary conditions are thus and so then later conditions will be such and such.

Unfortunately for us, many different conditions in the past, the elsewhere, and the undetectably small are compatible with our human situation and its observed development. Those different compatible conditions give us -- if we add no special assumptions or wishful thinking -- many different predictions.

Hence those beautiful laws, together with our present knowledge, leave us in a predictive vacuum.

Newtonian science allows for the logical possibility of conditions evolving as we know them, up till now, and then becoming radically unstable, from our point of view, at every level. It may be objected that I'm taking the usual skeptics' or empiricists' recourse to weird logical possibilities here. Surely we can discount those nightmares? They are incredibly improbable!

Fine, I will personally assent to this probability judgment -- but *that probability* is not supplied by Newton's science. There is no probability measure on the initial conditions which comes with that scientific theory. Well, then, can't we infer backward to initial conditions from known conditions today? Agreed that we can't do it by mere deduction, but what about probability? Bayes' Theorem is designed to address this question. Unfortunately Bayes' Theorem helps only those who have helped themselves to a prior probability. Given the evidence it will tell us how to arrive at posterior odds by multiplying the prior odds (between hypotheses) with the 'Bayes factor'. But we have to supply those prior odds ourselves. Science does not do so.[14]

I've focused on Newtonian science for concreteness, but the situation is no better today. Quantum mechanics gives us only *conditional* probabilities. If certain conditions obtain, then specified outcomes have calculable probabilities. But there is no probability for those initial conditions.

There are indeed attempts to replace or supplement quantum mechanics with true 'tychism' -- i.e. a theory which derives the conditional probabilities from a postulated absolute probability distribution. Examples are Bohm's hidden variable theory as recently elaborated, and the Ghirardi-Rimini-Weber theory.

Well, that is not science today. The Moorean Scientific Realist should be the last to base his arguments on anything but accepted current science. Anyway, imagine the news that new interpretations of quantum mechanics, or new developments in atomic physics, may save Induction after all! Little did the Newtonians know

2.3.3 What about the Special Sciences?

It is easy to imagine a Moorean Scientific Realist's reaction to my demurrals. True, s/he will say, science corrects and criticizes our inductive habits when they take certain concrete forms. But the important word is "corrects": the unreliable practice is replaced by a more reliable one. Thus it is a true advance when we stop 'predicting' the weather more than five days ahead: we are better predictors when we make fewer unreliable predictions. Meanwhile the same science has given us more reliable predictions for the next four days.

Granted, the 'basic' or 'fundamental' sciences do not do this for us. The Theory of Everything, or what passes for it in a given age, provides only a framework for the special sciences that provide the humanly usable information needed to live on this small planet, in this solar system, within this cosmic era. Hence the limitations of the 'basic' science are simply beside the point. The task for the Moorean epistemologist is to show the reliability of the methods actually used by science, under the historically given conditions, on the basis of what we have found out about those conditions.

Surely this is plausible, and commendably modest as well? But this plausible claim is not being made about the methods used to predict the weather, build bridges, or improve our agriculture. It is being made about the methods putatively making up the scientific method, the means by which the basic and special sciences are constructed. Those special sciences are biology, medicine, materials science,

psychology, Will these special sciences imply then that methods used to develop scientific cosmology, theories that tell us about reality *far beyond our small spatio-temporal region*, are reliable *when used by people living in that region*? Or will they be able to tell us only that very specialized inductions, say concerning our planet's atmosphere and global warming *within our own historic era* are reliable?

As concrete example, let us look at a very successful class predictions that we all know. For most familiar sorts of interaction we predict that entropy will increase, not decrease. We predict that the cream stirred into the coffee will not separate out again, either spontaneously or with further stirring, and so forth. We might call this the "thermodynamics of everyday life".

An epistemologist might well claim that these rules of prediction were initially arrived at by straight induction or abduction from pervasive features of the experienced world. But today, long after they became common practice, science justifies these rules, by implying that they are generally reliable. These small rules of prediction can be justified directly on the basis of thermodynamics, with its famous second law that entropy tends to increase and never decreases. (The "never" can more realistically be put in statistical form, giving us "moral certainty" that we won't encounter exceptions.) This is a good case to examine, for the second law of thermodynamics, even in statistical formulation, does not follow from basic science. It cannot be derived on the basis of classical mechanics, unless we add an assumption about initial or boundary conditions, whether outright or probabilified. It has to be regarded as a principle of one of the special sciences, applied chemistry perhaps, or applied physics.

If we now take for granted that thermodynamics is by now part of our knowledge of nature, then we know that the rules of the thermodynamics are reliable. Note well that they have then become simply applications of deduction and the statistical syllogism -- they are no longer inductions. Induction is not needed once we can derive our predictions from a special science which has become part of our knowledge.

Could we perform the same trick for the methods of induction, abduction, and the like which are assumedly used to construct all the sciences? A precisely similar justification would follow from our knowledge only if it included a sort of "universal second law", some knowledge which entails that those methods are reliable. If we had that, we would of course no longer need non-deductive, ampliative methods -- we would have replaced them with deductions from that "law". We have quite clearly arrived here at one of the famously impossible ideas in the history of induction: its justification on the basis of a principle of "uniformity of nature". But worse: we have very idiosyncratically arrived at the idea that this principle could be established by the special sciences, which investigate the particular conditions of human existence!

2.3.4 So, Scientifically Speaking, How Do We Reach Conclusions about Initial and Boundary Conditions?

The short answer is: not by any method whose reliability is implied by science.

Let me rephrase the question as follows: how does science proceed from basic science to the special sciences? Keeping his basic laws of mechanics as foundation, Newton adds the law of gravitation to describe this universe, he adds that there is one sun and six planets to describe our solar system, he adds that there is one moon to describe the more immediate gravitational forces on our planet Earth. Seen

in this way, the extension from basic to special science is simply by postulation. But of course that is not a realistic picture: Newton had reasons to choose the specific postulates -- he knew quite a lot about the phenomena and wanted to 'fit them into' his model.

This is the precise point where we notice a very deep assumption at work, which was again highlighted by Reichenbach.[15] In this 'fitting in' a certain assumption is made. I'll give it a special name:

Approximation Principle: if certain conditions follow from the ideal case, then approximately those conditions will follow from an approximation to the ideal case.

Newton demonstrates that from a very idealized, simplified description of the solar system, something approximating the known phenomena follows. Well, so what? What's the point of deriving true conclusions from a false premise?

The point is that this demonstration 'fits' the phenomena 'into' the Newtonian model if we assume the above Approximation Principle. On that assumption, Newtonian science has then successfully entered the domain of the special science and can take it over: provide it with new foundations, and both correct and extend scientific opinion in that domain in hopefully fruitful ways.

But is this Principle true? Newtonian mechanics does not imply that it is. It was well known long before the advent of Chaos theory that classical physics allows for cases in which that is just not so. Therefore the basic science cannot justify use of this Principle. But the special sciences, conceived of as extensions of basic science, are established under the assumption of that Principle -- indeed, that Principle is the only and the crucial link between the special sciences and their more basic parent science(s). Therefore the special sciences cannot justify this principle for us without circularity.

Science can certainly tell us that under certain conditions an ampliative rule we use *is* reliable. That is correct. For example, science may spell out conditions under which the forecast rule

If it rains today then probably it will rain tomorrow

may enjoy a certain degree of calibration. But then the question is whether those conditions obtain. And there we go again. Fortunately for us, we have put our trust in many such rules of thumb that had survival value, and have been nimble enough to discard ones which led us astray before they did mortal damage. Lucky for us!

2.3.5 Where Science Meets Prejudice: in Philosophy

But now you want to turn the tables on me and say: how did we arrive at that scientific knowledge we have in the first place? Assume Newton's Theory is true: how did Newton arrive at this truth? Assume our present science is true: how did our scientists arrive at his truth? They must have used induction and abduction. Perhaps the truths they arrived at do not suffice to show or demonstrate that this road by which we came to them was a good road to choose. But we were vindicated, and there must be a reason for it. So if science does not yet support our inductive methods and practices then we must just pursue science and extend it until this domain too is covered, or else add Induction and/or Abduction to the very foundation of all science.

Such a response is full of assumptions that should now be in doubt. If science does not support those ampliative rules, then it certainly does not underwrite the conviction that truths we know now were discovered by such means. Those ideas of Induction and Abduction, born and nurtured in the philosopher's armchair, are part of a pre-conceived view of what science *must* be like. The conviction

that science proceeds by such a method, and that therefore the burning question is whether we can have good, non-circular reasons to believe in their general reliability, was part of both the First and Middle Way. It has only led us into culs-de-sac.

Science alone does not imply the reliability of any humanly usable ampliative rule. To conclude such reliability we must *add to science*. We must add our own assumptions about our good fortune, or our own probability measure. Neither addition can be derived or justified independently, whether on the basis of science or on *a priori* grounds.

Let me make one thing very clear. I do not classify as irrational anyone who puts his trust in any properly formulated ampliative practice, whether it be inductive, abductive, or what have you.[16] But it is an illusion to regard confidence in such a practice as more than trust or faith that actual regularities are being caught in its noose.

At this point it would surely be pure irony to suggest that, to account for the rationality both of our trust in science *and* of trust in those special assumptions or prior probabilities, we must accept the idea of induction nevertheless. Apart from irony, I have only one conjecture: that the idea of induction remains alive because we *seem* to have no rival alternative. But in reality we do have an alternative; there is a Third Way; and to this we now turn.

III. Luck, Courage, and Technique

Reichenbach's history of empiricism as philosophical disaster extends down to the present time, along the same damning lines. Empiricists accepted tacitly both central assumptions and criteria of adequacy for epistemology from their predecessors. Despite continual weakening of this troubled heritage, they found themselves facing impossible tasks. Many fairy stories have this exact plot structure, but there the heroes and heroines really do end up doing the impossible! But perhaps, as Reichenbach tried to do, we can benefit from this diagnostic history to reconceive the challenge.

3.1 The New Challenge

Carnap suggested that in epistemology we must, in effect, design a robot who will transform data into an accurate probabilified description of his environment. That is a telling image for modern ambitions; it also carries precisely the old presuppositions which we have found wanting. We can already design pretty good robots, so the project looks initially plausible. But the project's pretensions to universality and self-justification send us out of the feasible back into the logically impossible -- back into the Illusions of Reason, to use the Kantian phrase.

Let's not design Carnap's robot. Instead, let us try to imagine the tragic protagonist who hopes to survive, flourish, and even enjoy these fleeting moments in existence, without illusions of security -- nor, certainly, any guarantee of a happy ending even for science in the ideal long run.

What will be the point of this exercise of the imagination? It has two points. The first is the point of all epistemology, past, present, and to come: to make sense to ourselves of our epistemic life. The second is that as philosophers we are caught in a historical dialogue. Surely all that admirable, astonishingly clever work carried on in service of the Illusions of Reason is not without value -- surely we can place it in a light that shows the true worth?

There exists already a great deal of work that we can draw on for an alternative, non-traditional

epistemology. I have described that elsewhere (see my (1989), Chs. 7 and 13), and I do not want to add to it here. Instead of focusing on the details, I want to convey its spirit. Given that traditional epistemology embodies false hopes never to be satisfied, we must try to find a different view of our epistemic condition, with new hopes and new dreams of its own.

3.2 What Does it Take?

At the far end of infinity a coin is spun, it will come down heads or tails. How will you wager? Reason cannot you choose either, reason cannot prove either wrong."

Blaise Pascal, *Pensées*

So here is our tragic protagonist, thrown into a world she never made, and she asks us: *What does it take?*

We answer her: it takes luck, courage, and technique; but the greatest of these is luck.

We do have something more to tell her, though, for each of these. The part on luck and courage will be the part the Pragmatists (and Existentialists?) contributed to that historical dialogue. The part on technique will astonish her. No amount of luck and courage, practically, will save her if technique be lacking -- and yet, technique, of which even a little will carry her a very long way, is empty.

3.3 The Parts of Fortune and Courage

If our pursuit of knowledge, however broadly or feebly construed, is to be successful, we must be *lucky* -- we have no way to constrain such fortune. This is the verdict on modern philosophy's misguided search for security.

The history of Earth has seen great disasters that spelled the extinction of almost all life, including the dominant, best adapted among species. In each case, *some* forms of life happened to be suited to the radically and suddenly transformed circumstances -- thus evolution on Earth continued after all. See who was lucky and who was not! Look to those survivors, they weave not; neither do they spin; but fortune smiles on them.

This is an insight into our condition; it is not skepticism. The insight leads to skepticism only when courage is lacking. (See further my (1988).) It takes courage to rely on the reasoned opinion and skills you have developed, given the contingency of their conditions of applicability.

I obviously need to say something more to link this remark to a serious discussion of epistemology. Part of my assertion is a diagnosis of philosophical skepticism and its pathology (cf. my 1988); I won't repeat that here. The other part is the Pragmatist theme that epistemology cannot proceed in isolation from value theory -- that the epistemic enterprise cannot be adequately conceived without attention to the role of value judgements, intentions, commitments, decisions, and other aspects of praxis.

The Pragmatist theme is the theme of courage -- a virtue which can enter when purely epistemic virtues are not far reaching enough. This theme has often been misunderstood. Of course courage does not by itself increase your chances of being right! There are heroes on all sides, winners and losers alike.

William James made a serious misstep when he presented this theme in his "The Will to Believe". It is not just the reader's fault if James reads as if he is presaging Norman Vincent Peale's "The Power of Positive Thinking". He gave the example of a mountain climber who comes to a chasm cutting through his route. To jump is to risk death by falling short and not to jump is to risk death by exposure. James

points out that the belief that one can jump is an almost necessary condition for jumping successfully. Thus he advises the mountaineer to believe that he can do it -- as if one could form a belief at will or for such ulterior motives.[17]

All of this is beside the point, and the rest of James' writing attests to it. The contrast between the mountaineer who jumps, perhaps to survive and perhaps to die, and the one who hunkers down till dawn, perhaps to survive and perhaps to die by exposure, is a contrast between two forms of life. The question is: which life does he find worth living? Once we see that there is no method, no rule, which will produce a rationally compelling conclusion, that is the question which comes to the fore. Virtues other than the ability to calculate profit and risk come to light.

We can certainly not identify courage here with the course of lesser prudence or of precipitate physical action. The illustrative example is merely illustrative, and illustrations oversimplify. Either course can be followed courageously; the courage does not lie in the action, but in how life is lived. The point is Pascal's: once reason comes to an end, once calculation is done and can carry us no further, other virtues come into play.

The point about non-epistemic values entering epistemology pertains not only to the subject (epistemic life) we are analyzing, but also to our analysis. Although I have stated the matter as neutrally as I could, value judgments are intruding not just on the individual's epistemic life, but on our epistemology. Only courage can take us out of skeptical despair, but who says that a life of skeptical despair is not worth living? We say so, and modern philosophy already said so in setting its goals for epistemology. But this consensus too comes from us, not from anywhere else.

Thirdly and finally, *while indispensable*, courage plays a better but smaller part if there is less need for raw or blind courage. That is so exactly when we have a better understanding of our situation.

3.4 The Reach and Limit of Technique

In its main use, *technique* denotes rather narrowly the skills acquired under special circumstances for special ends. These include inductive practices, like the soldier and sailor's separation of relevant from irrelevant factors for forecast and decision. In this narrow sense all technique is bound to its originating circumstances; and is defeated by change. Well-developed old skills pale before gunpowder and steam propulsion.

What is characteristic of modern science is the development of technique progressively less bound to special circumstances. We do not build houses suited to an ice age. But should a new ice age rapidly overtake us, we are confident that we have the general skills to modify presently applied skills, to adjust. The more theoretical our science, the less bound it is to special assumptions about our circumstances.

In the limit, *what is* technique not bound to any special built-in assumptions or historical constraints? It is logic and mathematics. In fact, the ultimate technique, and at the same time the fall-back position for all technique, is simply and exactly that: pure logic and mathematics.

What I would like to do in this last section is to make plausible two contentions. The first is that there is only so much science, or method, in scientific method as there is logic and mathematics. The second is that we can generalize this point to all epistemology, to conclude that rationality is in some good sense empty. Any truly coherent opinion is truly rational -- but this conclusion does not imply the dreaded 'anything goes' anymore than the skeptical 'nothing goes'.

3.4.1 The only method: logic and mathematics

So my first contention is that there is only so much science, or method, in scientific method as there is logic and mathematics. In one form this contention is a familiar one. For it is part of the orthodox Bayesian position that theory choice in science really takes the form of conditionalization -- that is to say, merely logical updating. Some Bayesians have provided elaborate historical reconstructions to support this contention. That is not the form in which I want to advocate it. As far as I can see, Bayesian retrospective reconstructions of scientific successes are pseudo-science, since they are always possible, and exclude nothing.[18] So this is not what I mean.

But the main point stands. All those successes of science which so many people have thought must have been produced by induction or abduction were actually successes of logic and mathematics. For they were initially good guesses under fortunate circumstances. But they were made effective by means of the precise formulation and disciplined teasing out of their implications through logic and mathematics. Newton's laws of motion, Darwin's theory of evolution, Einstein's recasting of kinematics would all today be on a par with Thales', Anaximander's, and Anaximenes' cosmologies if it weren't for that difference in how they were mathematically developed.

If this does not seem plausible, it is largely because the power of logical thinking tends to be discounted and unappreciated. It is also because of a myopic preoccupation with the felt need for things we cannot have, which diminishes the value of what we can have. One striking illustration of this is found in Thomas Pynchon's novel *Gravity's Rainbow*.

"A screaming comes across the sky. It has happened before, but there is nothing to compare it to now. It is too late. The Evacuation still proceeds...."

This is the beginning of *Gravity's Rainbow*, the description of a city at war. 1944: London is under attack by V-weapons. A grey stone house just off Grosvenor Square is ACHTUNG: Allied Clearing House, Technical Units, Northern Germany, one of a number of technical intelligence units engaged in evaluation and research. All are trying to find some guidance for deployment of the overwhelmed fire brigades, emergency crews, ambulances.

Each morning someone in Civil Defense publishes a list of yesterday's hits. Only one person appears to have any success, a statistician, Roger Mexico. "His little bureau is dominated now by a glimmering map,...an ink ghost of London, ruled off into 576 squares" -- a grid of 24 x 24 squares each representing a quarter square kilometer. Beside the map hang various diagrams: his curves and projections, based on the Poisson equation.

To the statistician Mexico "belongs the domain between zero and one -- the probabilities. A chance of 0.37 that, by the time he stops his count, a given square on his map will have suffered only one hit, 0.17 that it will suffer two,...." From these numbers Pynchon's reader can deduce (by means of Poisson's Equation) that Mexico's current predictions are for a count of approximately 540 missiles.

Every day he plots the new data. Every day, a better fit. "The rockets are distributing about London just as Poisson's equation predicts. As the data keep coming in, Roger looks more and more like a prophet. Psi Section people stare after him in the hallways...."

But the hopes this raises are vain. The statistical equation graphs the angels' view -- who see the invisible city depicted so accurately in this grid, these numbers, these curves. "Why is your equation only for angels" ask Jessica Swanlake, "Why can't we do something.... Couldn't there be an equation for us

too, to help us find a safer place?" But no, the very assumption on which Mexico's projection is based, is that there is no systematic clustering, no predictable pattern, only a constant mean density, no place better for hiding than any other. This is the paradox of the success that lies in the recognition of limits -- when reason replaces the hopes of magic and precognition. This is the predictable shape of chaos.

The fact is of course that Roger Mexico's findings are valuable not just for angels. For the city can organize its civil defense resources for maximal flexibility, so as to optimize response under conditions of minimal uncertainty. If psychics or inductive logicians could foretell the future, then! yes, then! But those are fictions. The mistake is to think that without those saviors we are helpless, when in fact we are in a position to do still infinitely more than we have been doing to help ourselves.

Philosophers, scientists, and mathematicians who tackled the detailed problems of formulating inductive methods -- they achieved a great deal. Their real achievement was to develop highly sophisticated statistical methods -- and what was true in all they said was what they demonstrated mathematically. What was lucky was that our environment was kind in to us important respects: convergence in infinite long runs turned out, in many important cases, to be reflected in humanly speaking large numbers -- often not even very large, even to our human eyes. There was in all this, to use Kant's words, only so much science as there was mathematics -- and besides science there was, happily for us, art, enterprise, and good fortune as well.

3.4.2 Not a road to skeptical disaster

Blindness and insight go together. The analytic tradition and positivist movements forced philosophy to come to terms with logic and mathematics -- but also taught us to dismiss them as empty. As soon as I suggest that all we have, and all it takes, in this world is luck, courage, and pure mathematics, you suspect me of skepticism again. For pure mathematics is only logic, and logic is empty, -- isn't that so? But I say, this empty logic carries us so far that it offsets infinitely more bad luck than the accidentally reliable inductive methods that were supposed to drive science forward.

Turning now finally to my second contention, let us examine the implications for rationality. I take it that what is rational is precisely what is rationally permitted. Thus we are rational in believing something exactly when we are not rationally compelled to believe the opposite. This implies, tautologically, that nothing more than staying within the bounds of reason is needed for this status of rationality -- not good reasons, not a rationale, not support of any special sort, not a pedigree of inductive reasoning or confirmation, nothing is needed above and beyond coherence. Thus any truly coherent position is rational.

Now this may strike you as absurd, because it seems to imply that 'anything goes', the ugly brother of skepticism's 'nothing goes'.^[19] But that is simply not so. If it were so, I would have come to an impasse and would have to admit that the constraints of rationality are more than coherence.^[20] But it is not so, and we have not in fact come to an impasse.

For at any given time, an individual responds to his experience in terms of and on the basis of his prior opinion and understanding (much of which is built into the language in which he frames that response). This means that although there may be leeway and options, they are highly constrained by the way in which we come to the new experience or new evidence. But the constraints are not transparently clear to us. It takes both salient new anomalies and highly technical sustained logical and mathematical investigations to reveal them -- and then perhaps to surpass in some way.

Logically speaking there may be infinitely many coherent alternatives to our posterior opinion, but almost none of those will be live options for us. This is not an accidental feature of our situation, but integral to it, exactly because we have prior opinion and prior understanding, which is not laid out for us on the side, like a textbook, but which is already our own. Any small part of it we can think of changing as part of our response, and any large part of it we can think as possibly changed, to any imaginable alternative, over the long run, in the course of time. But we do not have the wherewithal to respond in more than a limited variety of ways at any given time. Many of the imaginable alternatives to what seems plausible to us cannot be incorporated by us. At least, they cannot be so incorporated in a short span of time, through a coherent variation on our prior opinion, while doing justice to our new experience. So those alternatives do not go.

For example, cold fusion was imaginable, and to this day, it has not been logically ruled out, as far as I know. But no one has been able to include it in a hypothesis consistent both with the very large amount of background science in this domain and with the experimental results of the various laboratories which investigated the described conditions. So cold fusion is at this point not a live option for us. Nor does anyone know whether it is part of some logically possible physics that could have been as adequate as ours if developed by us under actual terrestrial conditions -- if it is, the angels may smile, but for us the point is purely academic.

It is important not to mis-gauge this point. It does not mean that there are constraints on rationality beyond coherence, and it does not mean that we are following ampliative rules after all. It means simply that the Neurath 'mariners repairing their boat at sea' form of relativism does not lead us into a damagingly enfeebled epistemic enterprise. The 'already on a boat' part defeats the skeptical argument based on the 'at sea' part.

So here is my conclusion. We supply our own opinion, with nothing to ground it, and no method to give us an extra source of knowledge. Only the 'empty' techniques of logic and pure math are available either to refine and improve or expose the defects of this opinion. That is the human condition. But it is enough.

ENDNOTES

1. This is the standard 'textbook' story; I do not consider it historically adequate, but its purpose here is to highlight certain alternatives in epistemology for the empiricist tradition to which Reichenbach (as I do) wished to adhere.

2. Here is the entire passage: "[I]n developing his own philosophy, the empiricist unconsciously accepted the fundamental thesis of rationalism, according to which genuine knowledge has to be as reliable as mathematical knowledge, and thus was pushed into the hopeless position of proving that empirical knowledge was as good as mathematical knowledge. Historically speaking, this rationalist commitment of empiricism found its expression in the fact that empiricists were always on the defense On the other hand, those empiricists who were too honest to deceive themselves had to end up as skeptics. The empiricist skeptic is the philosopher who has freed himself from the intoxication of rationalism but still feels committed to its challenge; who regards empiricism as a failure because it cannot attain the very aim

which was set up by rationalism, the aim of an absolutely reliable knowledge. [...] That is true of the most prominent of the empiricist skeptics, of David Hume, whose philosophy is both the first consistent construction of empiricism and its breakdown." (op. cit., 139-140)

3. A common interpretation of Peirce, of late forcefully contested by Elisabeth Lloyd, has him cutting the Gordian knot by equating truth with what we will believe in the long run. This presupposes that the methods of science are such as to guarantee convergence to stable opinion -- something he could not demonstrate. It also still leaves open the problem of why we should think that our 'corrections' (when not merely weakenings) will typically or often be improvements (except in the sense that in a divine retrospective they were stages on the way to what is stably believed in the long run).

4. See A.P. Dawid (1985) which also includes a summary of the preceding discussion by Box, Dawid, Lindley, Oakley and others.

5. If there is underdetermination, it will pertain to the way in which the forecaster draws on background theories and assumptions -- or on his own talents and instincts -- but the distinct forecast probabilities will converge uniquely. The existence of a computably calibrated rule under favorable conditions can be shown on the basis of the following. There are only denumerably many specifiable classes to be considered, since each challenge must be given by means of a computable function. Now if the 'real probability' in nature has a domain including these classes, then there will be a possible sequence in which the relative frequency of each such class matches its 'real probability'; indeed the set of such 'good' sequences has measure 1 ('real probability' 1).

6. The first to see the connection between this idea and that of a universal algorithm was, to my knowledge, Hilary Putnam (1963, 1985). In his contribution to the 1963 volume on Carnap's philosophy he gave a diagonal argument to refute the hope for a 'universal' inductive method. See further Gaifman and Snir (1982), and Oakes (1985) with the discussion between Oakes, Dawid, and Schervish accompanying that article on the next three pages. The argument is related to the simple proof that there is no 'universal' computable function. For suppose the computable functions are $\{F(i)\}$ and $U(i,n) = F(i)(n)$ for all i, n . Then $U+1$ is computable if U is, and differs from each $F(i)$ at argument i ; hence is not computable. Of course this shows that the list $\{F(i)\}$ is not effectively specifiable. The same is true of the computable forecast systems and seems likely to be true of any significant subclass we might want to consider (such as those basing their calculations on 'acceptable' scientific theories plus data, etc.).

7. Because the matter is abstract, the criteria applied may seem unrealistic. But the results are rather robust. Consider for example the 'swamping' results which suggest that the evidence will "swamp" a forecaster's initial probabilities, through conditionalization, and so bring them into closer and closer agreement with the 'real' probability. Like Reichenbach's vindication of the straight rule, these results seem to be interpreted with blinders on. In Oakes 'diagonal' procedure, we can take the list of forecasters to be ones who all conditionalize, but start with a different prior. Alternatively, suppose that we know or assume that the 'real probability' function belongs to a certain denumerable set FF . Then we can choose a mixture (convex combination) of the members of FF as our U_r -prior. Conditionalization on a sufficiently rich sequence of data will, with 'real probability' 1 transform that prior into a function that converges to that 'real probability'. The problem, as Dawid points out, is that this U_r -prior is specifiable only if the

class FF has an effective enumeration. Thus our initial assumption or knowledge claim will be insufficient or incredible.

8. In many discussions of the adequacy of statistical methods the picture has undoubtedly been confused by demonstrably rational *subjective* certainty of success for users of statistical methods. If someone believes a statistical hypothesis which assigns probabilities to classes of natural events -- be they coin tosses or subatomic processes -- that believer can and must be 100% certain that the hypothesis will be vindicated if the trials be fair and the run long enough. The theorems concerning 'laws of large numbers' transpose to computable forecast systems. But that demonstrated certainty is predicated on the supposition that we have latched on to the right rule -- it is certainty of success for the believer, and really just another aspect of coherence.

9. As I indicated before, we can also see Reichenbach as ambiguous or vacillating between the option we have just examined, and the one to which we are about to turn -- or we can think of him as progressing from the former to the latter, while remaining to some extent in the grip of the older ideal. See for example p.77 of his (1959), translated from an article he published in Erkenntnis 1930.

10. In this connection see also the literature centering on Laudan's naturalistic approach to epistemology, reviewed in Schmaus(1996). See also an article by Stephen Leeds, which really conveys the idea quite well, if a little obliquely: "all that matters is the fit between our methodology and the way the world actually is. If the sorts of accounts we value most highly are in fact the way the world is constructed -- if the laws of Nature involve geometry, conservation principles, atomistic reduction, and causal influences, say -- then our methodology has a fairly high chance of selecting truths." (Leeds, 1994; p. 203).

"[T]he Realist ... will tell a story according to which the world in some sense has guided us to the right standards of explanation and simplicity. [...] A careful Realist will present his story as a corollary to what he takes himself already to know about the world, and in particular, a consequence of the fact -- remarkable, but solidly integral to the Realist's world picture -- that some, although by no means all, of the most fundamental laws of Nature involve processes which appear also on the macroscopic level It is for this reason that, a Realist can argue, a creature who can find his way around the macroscopic world must inevitably be equipped with the 'right' standards of simplicity and explanation. (*ibid.* pp. 203-4)

11. Hailey and Helfand (1996) of Columbia University note the following concerning the crash of T.W.A. flight 800: "Approximately 3,000 meteors a day with the requisite mass strike Earth. There are 50,000 commercial airline takeoffs a day worldwide. Adopting an average flight time of two hours, his translates to more than 3,500 planes in the air [at any moment]; these cover approximately two-billionth of the Earth's surface. Multiplying this by the number of meteors per day and the length of modern air travel [30 years of high volume air-travel] leads to a 1-in-10 chance that a commercial flight would have been knocked from the sky by meteoric impact."

12. Casti (1990), p. 99: "As a specific example ..., let's consider the British Meteorological Office model, which is divided into fifteen atmospheric layers starting at the surface of the Earth and rising to a height of twenty-five kilometers In this model, each level (layer) is divided into a network of points about 150 kilometers apart, giving around 350, 000 grid points. Each of these points is assigned values of

temperature, pressure, wind, and humidity every twelve hours from observations taken over the entire globe."

pp. 100-101: In Figure 2.9 we see the correlation coefficient between the predicted and observed values of the changes in height of the 1,000 millibar (mb) level for forecasts made by the British Meteorological Office ... during the period 1968-1984. [...] By [the root mean square] criterion, in 1974 the error of the 3-day forecast was only 20 percent less than the persistence error, which is the error made by predicting that the weather tomorrow will be the same tomorrow as it was today. [...] But by 1984 the corresponding figure was 52 percent, and the degradation to the 20-percent level did not take place until day 6 of the forecast."

p. 127: "When it comes to weather forecasting, the most common reason for predictions to diverge is differences in the analyses of the raw data that goes into making up the initial conditions. These errors in the initial conditions, coupled with the intrinsic sensitivity to initial conditions of the dynamical equations, are the primary obstacles Studies have been carried out to test long-range predictability using the so-called perfect model assumption, in which the mathematical model of the atmospheric dynamics is taken to be an exact representation of the actual physical process itself. These studies show that errors reach their asymptotic, or saturation, level after a period of 14 to 20 days the error growth rate during the early stages has a doubling time of around two days."

13. When Goodman divided predicates into projectible and not projectible, he was simply naming the problem. If projectibility classifies the cases in which induction is successful, then projectibility is a function of parameters to which induction has no access. If it classifies the cases in which we have the habit or find ourselves strongly inclined to generalize or extrapolate, we'll just be plumb lucky if it has much overlap with success.

14. I'm assuming the failure of the long search, starting already early in the 18th century, for a way to derive prior odds by symmetry arguments. See my (1989), Ch. 12 and recent discussions such as Albert (forthcoming), Craig Callender (forthcoming).

15. A preview of chaos theory -- I refer here to Reichenbach's striking and little appreciated demonstration that indeterminism was all along consistent with the success of classical physics. He gives an example from statistical mechanics to show that the classical deterministic laws can, with appropriate initial conditions, lead to behavior that is, by measurement, indistinguishable from genuine indeterminism. The example is exactly one in which what I call here the Approximation Principle fails. See his (1991), pp. 93-95.

16. I say "ampliative practice" here, rather than "ampliative rule" or "ampliative method". As I have argued elsewhere, to commit oneself to a rule of that sort is to land in epistemic incoherence. But it is possible to make ampliative revisions in one's opinion without following a rule that dictates how to do it, and this practice cannot be objected to in the same way. See my (1989), Ch. 7, sect. 4-6.

17. Even Pascal had not thought that one could, except perhaps through the long, slow option to *abêtir* one's mind through daily rituals. Yet who knows what the enterprising American divine could not do for us? Perhaps a training course in 'positive thinking' could produce a psychological surrogate for belief

with the same desirable effects!

18. Bayesian conditionalization is simply logical updating of opinion when new evidence must be assimilated. It is a good rule, and when applicable the only rule, for updating opinion under 'normal' circumstances, but progress would stumble to a halt if we were bound to it. The alternative to Bayesian bondage, however, is not embrace of ampliative rules -- even if we could formulate those satisfactorily. See further my (1989), Ch. 7.

19. I take it that this is Herman de Regt(1996)'s reason for suggesting that besides coherence, rationality requires that we have pragmatic reasons for belief. But as he points out, in effect, pragmatic reasons are from an epistemic point of view ulterior reasons. It is not possible (not pragmatically coherent!) to think that one believes something and believes it for any reasons that do not make it more likely to be true. (That is how I would put Bernard Williams' well known point against the possibility of believing at will.) My view is that all coherent states of opinion are rationally permissible, but that given one's extant state of opinion, very few changes (whether in response to experience or to deliberation) are live options. To this I would add that the concept of reasons for belief in traditional epistemology does not answer to the ordinary employment of such phrases on which it is ostensibly based. If I advance a view and you ask me for my reasons for holding it, I will try to select something that I think you also believe (or perhaps something that you, and not necessarily I believe) which will count as support for you. Thus conceived the concept of reason is limited to the context of dialogue, where it displays a certain relativity, and does not point to any kind of hierarchy in my own beliefs.

20. Notice that this would mean that a Bayesian agent (the pure conditionalizer described by the orthodox Bayesian) would not be rational except by accident, nl. if s/he had accidentally started from some point that satisfied the extra constraints and if those constraints were preserved under conditionalization on evidence received. Now the second (the preservation) is no more likely than the first, and not so in general for arbitrary evidence, if (as would seem) those additional constraints impose simplicity, explanatory power, and the like in some not purely logical sense.

BIBLIOGRAPHY

Albert, D. "On the foundations of quantum mechanics and the foundations of statistical mechanics", forthcoming in PSA 1996, vol. 2 [Supplement to Philosophy of Science].

Callender, C. "What is the problem of the direction of time"; forthcoming in PSA 1996, vol. 2 [Supplement to Philosophy of Science].

Casti, J. L. Searching for Certainty: What Scientists Can Know about the Future. (NY: Wm. Morrow, 1990).

Dawid, A. P. "Calibration based empirical probability" Annals of Statistics 12(1985), 1251-1273.

Dawid, A. P. "Comment [on Oakes (1985)]", Journal of the American Statistical Association 80(1985), 340-41.

de Regt, H. "The second best approach to the problem of scientific realism: rationality of belief" in Douven and Horsten (1996).

Douven, I. and Horsten, L. (eds.) Realism in the Sciences. Louvain: Leuven University Press, 1996.

Gaifman, H. and M. Snir, "Probabilities Over Rich Languages, Testing, and Randomness", Journal of Symbolic Logic 47(1982), 495-548.

Hailey, Ch. and Helfand, D. "In T.W.A. 800 Crash, Don't Discount Meteor", Letter to the Editor, New York Times Sept 19, 1996.

Laudan, L. "Peirce and the trivialization of the Self-Corrective Thesis", in his (1981), pp. 226-251.

Laudan, L. Science and Hypothesis: Historical Essays on Scientific Methodology (Dordrecht: Reidel, 1981).

Leeds, S. "Constructive empiricism", Synthese 101 (1994), 187-221.

Levi, I. "Induction as self-correcting according to Peirce", pp. 127-139 in Mellor (1980).

Mellor, D. H. Science, Belief, and Behavior. Cambridge: Cambridge University Press, 1980.

Oakes, D. "Self-calibrating priors do not exist", Journal of the American Statistical Association 80(1985), 339.

Peirce, C. S. Collected Papers, ed. Hartshorne, Weiss, et al. 8 vols. (Cambridge, MA: Harvard University Press, 1931-1958).

Putnam, H. "Reflexive Reflections", Erkenntnis 22 (1985), 143-154.

Putnam, H. "Degree of Confirmation and Inductive Logic", pp.761-783 in Schilpp (1963).

Reichenbach, H. The Direction of Time (Berkeley: University of California Press, 1956; new ed. with foreword by Hilary Putnam, 1991).

Reichenbach, H. Modern Philosophy of Science, New York 1959.

Reichenbach, H. "Rationalism and empiricism: an inquiry into the roots of philosophical error", Presidential Address to the American Philosophical Association, Pacific Division (1947); Philosophical Review 57 (1948); reprinted pp. 135-150 in his (1959).

Rorty, A. and B. McLaughlin (eds.) Perspectives on Self-Deception. (Berkeley: University of California Press, 1988).

Schervish, M. J. "Comment [on Oakes (1985)]", Journal of the American Statistical Association 80(1985), 341-342.

Schilpp, P. A. (ed.) The Philosophy of Rudolf Carnap (La Salle, Ill.: Open Court, 1963).

W. Schmaus, "The empirical character of methodological rules", PSA 1996, vol. 1 [Supplement to Philosophy of Science 63 (1996)], S98-S106.

van Fraassen, B. C. "The peculiar effects of love and desire", pp. 123-156 in Rorty and McLaughlin (1988)

van Fraassen, B. C. Laws and Symmetry. Oxford: Oxford University Press, 1989.