Philip Catton
University of Canterbury
Aotearoa / New Zealand

The Justification(*s*) of Induction(*s*)[1]

Induction is 'the glory of science and the scandal of philosophy'. I diagnose why. I call my diagnosis a "disappearance theory of induction": inductive inferences are not themselves arguments, but they synthesise manifold reasons, mostly tacit, that are. Yet the form of all these underlying arguments is not inductive at all, but rather deductive. Both in science and in the wider practical sphere, responsible people seek the most measured way to understand their situation. The most measured understanding possible is thick with arguments in support of every last belief. To achieve such an understanding is richly synthetic. Science has become systematically good at progressing towards this aim. But by virtue of their analytical orientation many philosophers are predisposed to misunderstand the nature of measurement, and thus to fall into confusion about the reasonableness of science. In considering an inductive inference, philosophers have expected to see one argument, rather than many; supposing that there is one argument, they have sought to describe its form; and then they have even attempted to establish a general kind of warrant for such a form of argument. Actually some significant philosophical contributions have issued from such work, but the worth, I argue, of these contributions, can only properly be appreciated when they are comprehended within the perspective that I defend. I also discuss how much more natural it is from the standpoint of synthetic philosophy (of the rationalists) rather than analytic philosophy (of the empiricists) to embrace the ideal of a most measured understanding, and in its light understand the integrity of both scientific and everyday beliefs.

§1. Introduction.

The foundational difficulty known as the problem of induction is simple to explain yet it apparently threatens almost the entire sweep of what we presume to be our knowledge. Faced with this ostensible

[1] I thank William Harper, John Norton and Stephen Sharp for helpful discussions.

problem we might hope to find some equally simple and easily stated solution to it—some single solid point around which to leverage our empirical thinking and thus re-establish our hold on our presumed knowledge of the world. I believe that to be gripped in the first place by the problem of induction, and to fashion such a hope as this for its solution, are two faces of a single coin. I also believe that this coin is counterfeit, good only for being recognised as such so that we may be better discerning in the future of what is an authentic philosophical conundrum and what is not.

David Hume impoverished people's understanding of what inductions are, and helped give false focus to the question of the warrantability of inductive inferences. After explaining and criticising these confusions of Hume's, I will trace their philosophical legacy into the twentieth century, the heyday of analytic philosophy. Within this analytic tradition, workers have attempted to characterise induction in general terms and thence to solve the supposedly general problem of induction. I will explain why this approach is both mistaken and doomed. In order both to find a better way forward, and indeed to create a perspective into which certain insights from analytic philosophers can best be incorporated, I argue that we must reappropriate the word 'induction', so as to understand that inductions are synthetically styled inferences that are not arguments at all. By discussing what it is to analyse the worth of such a synthetic inference, I reveal both the strengths of the analytical orientation in philosophy and also its limits.

I similarly reveal the strengths and the limits of the opposite, synthetic orientation in philosophy, epitomised in the philosophies of historical rationalists. In order to illuminate actual inductive inference making both in everyday life and in science, I show how important it is to win back some of the insights of that earlier tradition in philosophy, while at the same time championing some of the insights and the demand for clarity within the analytic tradition.

On the view that I criticise, inductive inferences are supposed to be arguments. It is then a philosophers' task to describe well the general form such arguments have, and to warrant our arguing in that way. Over against these persuasions I contend that inductive inferences are generally not arguments. To expect, falsely, that they are arguments, and thus that they have an explicit form, is to cross wires with the paradigm for deductive inference. I argue nonetheless that individual inductive inferences, in science and in everyday life, are often reasonable, and that if one is reasonable, then it will be so in ways we can analyse. Yet what such an analysis will discover to us is that inference's very own way of being reasonable, specific to its subject matter and the epistemic circumstances of those who wield it. The analysis in question succeeds in bringing to light the reasonableness of the inference only in so far as it illuminates underlying deductive features of the inference, that is to say, tacit arguments, deductive in form, in its support. There is no end to how rich these reasons may be.

In my picture, induction is to deduction much as coherence is to consistency. Synthetic philosophers emphasise how different the concept of coherence is from that of consistency and they are right to do so. Yet analyse coherence, and you reveal nothing but forms of consistency. The only difference is that you can never complete the analysis. (You also have to be synthetically discerning and creative in how you pursue the analysis.) There are forms of consistency but there is no form of coherence, since coherence is not the same as consistency, and yet when you analyse coherence you find nothing but consistency. Similarly, there are forms of deductive argument but there is no form of induction. And because there is no form of induction, there is also no question of warranting it.

According to my picture, we can no more eliminate the notion of induction in favour of that of deduction than we can eliminate the notion of coherence in favour of that of consistency. Yet that does not mean we can spell out a general notion of induction, or find within such a general notion anything at all to analyse. To find something to analyse, we need to focus in on a specific inductive inference. Then, if the inference in question is at all reasonable, there will be no end to what we can analyse. Yet the manifold significant reasons in its support will all prove under our analysis to be deductive in form.

If we analyse the reasonableness of inductive inferences that are reasonable, we discover manifold significant reasons in their support. But this does not in the least way suggest that inductions in general are reasonable, and of course it in any case is false that inductions in general are reasonable. The ostensible overarching issue about the justifiability of induction is really a will-o'-the-wisp.

The history of philosophy from the time of Hume to our own suggests a conclusion on the contrary that there is a terrible difficulty about saying what induction is and how it is warranted. I next begin my examination of that history, by considering the contentions of Hume.

§2. Hume on induction.

Hume first articulated his famous problem about induction in 1739.

- **Suppose that one hundred per cent, or alternatively some lower figure X per cent, of F's so far observed have had property G. Hume asked: have we on either account any right to form a definite expectation concerning the frequency in the future of F's being G? Hume concluded not: 'even after the observation of the frequent or constant conjunction of objects, we have no reason to draw any inference concerning any object beyond those of which we have had experience'** (*A Treatise of Human Nature*, ed. Selby-Bigge, p. 139).
- **Famously, his reasoning to this conclusion concerns two horns of a dilemma, corresponding to the two branches of "Hume's fork".** The reason, if there were one, to infer, from our past observation of a specific frequency of conjunction, any particular conclusion concerning the future, should, according to the fork, either be logical (that is, based upon the principle that whatever is such that its contrary implies a contradiction is true), or it should be empirical. But since no contradiction is implied by the idea that the frequency of F's being G might

change, the reason in question, if there were one, could not be logical. Yet, on the other hand, an *empirical* reason could only be circular, and thus come to nothing. For it could only amount to an inference from past such projections mostly faring all right to an expectation that future such projections will likewise mostly fare all right. The question would then become how *that* inference is reasonable, and an infinite regress would be under way.

Perhaps Hume's discussion is actually ironic, and as such it is in fact intended to be a reductio ad absurdum of the analytic orientation he overtly adopts. Yet that is not how his many followers have read him. So, in keeping with the traditional interpretation, I will treat the position just discussed as Hume's own.

It is important for us to take stock of what Hume has offered us, for almost all of it is in fact misplaced. Hume's contentions crystallise what is wrong with a purely analytic approach in philosophy to the question of how we learn from experience. I shall enumerate seven initial criticisms. In later sections I shall expand this to a wider list of reservations about Hume's overall stance.

(1) **Hume's problem concerns simple enumerative induction—a will-o'-the-wisp.** The form of inference that, according to Hume's argument, can never be reasonable, is that of simple enumerative induction. Hume's problem concerning induction is about licensing in general the following two inference forms: from 'All observed F's are G' (alone) to the categorical claim 'All F's are G', and from 'X per cent of observed F's are G' (alone) to the categorical claim 'X per cent of F's are G'. This problem is, however, wholly inconsequential, for simple enumerative induction is *never* used, either in everyday life or in science. We never infer from the observations alone; our epistemic situation is always rich with relevant collateral information and other already present theoretical beliefs. (In both the next section and several that follow, I will patiently illustrate why I say with such confidence that this is so.)

(2) **Yet Hume contends that his problem impugns almost the entire sweep of empirical knowledge.** Hume argues that unless simple enumerative inductive inference can be licensed, we are without good reason to augment our ways of thinking in any way beyond on the one hand the trifling truths of logic and on the other hand truths about specific empirical matters so far

observed. And as is well known, any number of very fine analytic philosophers have felt the force of Hume's concerns about this. Bertrand Russell, for example, admitted that without a solution, which he could see no way to provide, to Hume's problem of induction, he also could see no way to reason a man who thought himself a poached egg out of that persuasion. That is to say, Russell believed that Hume's problem impugns virtually the entire sweep of our presumed knowledge. Yet pace not only Hume but also the many analytic philosophers who have followed him, the claim that Hume makes here, about simple enumerative induction, is actually nonsense. That as a first step within our quest for contentful knowledge we would need to license simple enumerative induction is a claim so large and unreasonable that it is tantamount to insisting that there is a *necessary* condition upon our ever having contentful knowledge that is also a *sufficient* condition for our being outright insane. For it is completely straightforward to generate instances of simple enumerative inductions that no sane person would make. So the general licensing of this form of inference would itself signal insanity. Fortunately, Hume is mistaken. In order to have the right to claim to possess contentful knowledge we in fact would not need to license the form of inference to which Hume draws our attention. It is true that in order to have the right to claim to possess contentful knowledge various inductive inferences that we make each need in some way to be warranted. But none of these inferences is a simple enumerative induction, and the ways that any two of them are warranted need not be one and the same. I intend to reappropriate the word 'induction', which in my view was misappropriated by Hume, and apply it to a synthetic style of inference about which I say more below. At the same time I see Hume's problem as entirely irrelevant to the synthetically-styled inductions I say we use all the time.

(3) **Hume specifically considers only inferences to generalisations that are of the logically simplest form.** The upshot, according to Hume, of his problem concerning simple enumerative induction, is more specifically that *contingent generalisations* cannot be known. The contingent generalisations that Hume has in mind are logically utterly simple in form: All F's are G, or, X per cent of F's are G. Philosophers who attempt to solve or dissolve the ostensible general problem of induction typically accept that theoretical inference is primarily simple

enumerative induction to conclusions that are simply structured generalisations such as 'All F's are G' or 'X per cent of F's are G'. Much of the contemporary literature on laws of nature repeats this mistake, treating what laws are, or what laws are extensionally, as simple generalisations, when in fact the most illuminating laws are (or are extensionally) significantly richer than this from the standpoint of logic. Typical conclusions, law-like or otherwise, that people infer to inductively in everyday life or in science are logically much richer than 'All F's are G' or 'X per cent of F's are G'. I intend to illustrate shortly why I say this. Why it is important is the following.

(4) **Logically richly structured generalisations are however key to acts of measurement.** If theoretical contentions were all generalisations of the extremely simple forms 'All F's are G' or 'X per cent of F's are G' then it would be impossible even with theoretical contentions in tow to make non-trivial use of empirical facts in the deduction of other theoretical contentions. In short, *measurements* would be impossible. For in any measurement, our purpose is to deduce a theoretical conclusion from a phenomenon, in a way that employs a background of already theoretical further assumptions. I will discuss this in

detail later on. For the moment it is important to note how far Hume was from being the sort of thinker ever to have himself performed a scientific measurement. His unfamiliarity with experimental methodology shows itself to us in how he characterises as bare enumerative induction the supposed inference form for scientific theorising. Were our theoretical thinking never of the requisite logical richness for acts of measurement to be possible, we indeed would have to call induction merely a leap from the particular to the general, so that all inductions indeed would be simple enumerative inductions, and our situation would be every bit as hopeless as Hume contends that it is. But if our background assumptions include logically richer generalisations, then an actually demonstrative inference that uses empirical considerations to reach a theoretical conclusion does become possible. Thus, should I already believe that *if any F is G then they all are*—a generalisation that logically involves not one quantifier but two—and should I observe even one single F that is G, then it would be only *logical* for me to conclude that all F's are G. I call the use of any such form of inference an act of measurement. Taken on its own such an inference is, of course, deductive rather than inductive. But inductions in everyday life and in science are typically reliant on a nexus of more or less careful measurements. Whenever a view in everyday life or in science seems to us careful and well considered, we call it “measured”, and call it this precisely because it is. None of the measurement inferences that support it is completely cogent on its own of course, because of the fallibility of its background assumptions. I will amply illustrate this idea later; here, it is enough to observe how completely Hume leaves it out of account. For, among other things, Hume blocks our considering generalisations that are of the logical richness that is requisite for any act of measurement.

(5) **In light of measurement, Hume’s fork is a false dichotomy.** The inferences that people actually make show that Hume’s fork is a false dichotomy. And to say this is not simply tendentiously to invoke Immanuel Kant’s suggested tertium quid, the synthetic a priori. Or if it is, it is to render that idea down in part to an unexpectedly mundane consideration. Hume’s dichotomy is a false one because any induction apart from a facile simple enumerative one will involve broader reaches of our presumed knowledge. It will combine a variety of implicit or explicit acts of measurement, each of them fallible and uncertain because of

the fallible and uncertain theoretical assumption or assumptions that it employs. But the induction which combines these various implicit or explicit acts of measurement will be thus synthetic, and will prioritise various theoretical judgements that are already in place to the evaluation of the empirical evidence at hand. Hume is totally unprepared to acknowledge this kind of synthetic inference not only because of his fork, but also because acts of measurement would be impossible if every theoretical assumption had to have the simple logical structure 'All F's are G' or 'X per cent of F's are G'.

(6) **Hume fails to understand the logical form of measurement inferences. His discussion sets a trap for the unwary, so that whoever falls into it becomes either a skeptic like Hume or at least a hypothetico-deductivist. But the hypothetico-deductivists likewise fail to understand the logical form of measurement inferences. Their position starts as a concession to Hume and fates them to follow Hume into his conundrums about induction.** Hypothetico-deductivists have the grossly inadequate conception of the role of measurements in science, according to which measurements produce a merely elementary instance of a theory whose only logical function relatively to

theory is to *test* it. Some, the "critical" hypothetico-deductivists, such as Karl Popper, believe that passed tests in no way elevate the probability that a theory is true. Other, "inductivist" hypothetico-deductivists, believe that the successful passing of tests can confirm a theory, elevating the probability that it is true. Either way, hypothetico-deductivists insist that the logical connection between theory and evidence is from the theory (i.e. from the level of the general) to (particular) empirical predictions. They contend that there is no logical path from evidence (which is particular) to theory (which is general). They are thus quite evidently blind to the actual logical form of acts of measurement. For every measurement inference explores a logical path that is precisely from something empirical and particular to a general, theoretical, conclusion. Such an act of measurement of course depends as well on a host of theoretical background assumptions, for it is these that direct us to the salience for further theory of the empirical phenomenon in question. In order to appreciate the actual logical moment of individual acts of measurement one must not fall under the sway of hypothetico-deductivism. Two factors in philosophy of science pedagogy predispose philosophers to embrace hypothetico-deductivism, however. The first is the tendency to illustrate theories and theoretical laws using logically simple generalisations such as 'All F's are G' or 'X per cent of F's are G'. The second is that both Hume's fork and Hume's problem are easy to teach, and they seem to most teachers especially fetching examples of a philosophical insight. The easiest way to invite students to move on from these supposed insights of Hume is to suggest, falsely as it happens, that there is after all a logical function for evidence in relation to theory—but this function is purely critical, to test the theory. This suggestion is not only false, but in relation to the supposed problem of induction is also to no avail, as is sharply illustrated for us by the way that Popper's anti-inductivist philosophy cheats on itself and thereby fails. Yet it is for these simple reasons concerning standard pedagogy that philosophers of science generally convince themselves to be hypothetico-deductivists. To suppose with the hypothetico-deductivists that the only logical path is from theory to the evidence, is, however, in effect a way of following Hume into his errors.

(7) **Hume sets us to considering an illusory issue: how there could be any warrant for the first-ever theoretical inference.**

Hume might reply as follows to my insistence that a considered view will always be a measured one, and thus that measurement inferences are often used. Since such inferences themselves rely on theoretical assumptions, it remains to explain how there can be a warrant for those. One is chased by this consideration ultimately to the question how the first-ever theoretical inference can have been warranted; and it is of course impossible that the first-ever theoretical inference could have been made in a measurement, or measured, way. Thus the suggestion would be that until we explain how we might have warranted the first-ever theoretical contention, we have no way of explaining how we might have warranted any theoretical contention whatsoever. I will argue, on the contrary, that cognition is *always already* rich with theoretical contentions, and that the exercise of warranting a new contention depends on this being so. Thus the supposed cognitive task of explaining how we might have warranted the first-ever theoretical contention is ill conceived and actually irrelevant.

§3. How, in the face of Hume's problem, is science at all possible? Case study: the possibility of *geology* as a science.

Hume's problem supposedly confronts us with the heady general question how science is at all possible. Rather than asking first this heady question, let us instead consider for a moment how *geology* as a science is at all possible. The answer to that more specific question seems to me relatively straightforward to give. We can learn from this answer something significant about why Hume's singularly analytically styled investigation fates him to be a sceptic about induction and thence about the very possibility of our having any science at all. We can also appreciate the more clearly how inane this outcome is.

Quite without recourse to geological understandings we know enough to realise that what we call the earth is a material agglomeration. We are impressed by many reasons to anticipate that such a material agglomeration will be somehow or other patterned, at various spatial and temporal scales, in rich, nuanced, but also robustly identifiable and reidentifiable respects. Thanks, therefore, to various other spheres of knowledge that we have, and thus for ways we have of understanding matter, there are all sorts of initial geological questions we can form for ourselves about what might be the more or less stable behaviours and structures of the earth as material agglomeration, and about the techniques we can devise for empirically ascertaining the answers to those questions. The other sciences we develop, including most especially chemistry, physics, and life science, help shape and enrich the geological questions we ask and the techniques we can devise for empirically pegging down the answers to those questions. How is geology at all possible as a science? Well, I have just sketched the answer.

To add detail to this answer let us investigate how a particular geologist can warrant a specific theoretical contention. Let's consider Alfred Wegener and his celebrated contention that the continents have moved. Wegener in fact adduced singularly impressive evidence that the continents have moved. That the continents have moved is, of course, a theoretical contention if ever there was one. What was it for Wegener to offer empirical evidence in support of this contention?

Well, one thing Wegener did was to examine actual terrestrial landscapes. Of course, he did so using his considerable practical knowledge as a geologist. In Africa, he recognised bits of old moraines. That is, he recognised certain extant features of a landscape as expressing the past action of glaciers that had once existed but had later long ago receded and disappeared. It is within the practical wherewithal of most every geologist to be able to recognise a moraine as a moraine. Wegener went further however, and carefully analysed the composition of these bits of old moraine. That is, he quantitatively assessed the proportions of various types of rock, the incidence of various kinds of fossils, and so on. In South America, Wegener likewise recognised bits of old moraines. Again he analysed their composition. Astoundingly, he found that some bits of old moraines in Africa were utterly alike in constitution to some bits of old moraines in South America. They had identical proportions of various kinds of rock, identical incidences of various kinds of fossils in those rocks, and so on. But as a geologist Wegener knew to expect no two moraines to be alike in constitution unless they were pushed up by the same glacier. Geological variety, Wegener well knew, makes it utterly unlikely that two moraines, pushed up by two different glaciers in two different places, should end up indiscernible from one another in constitution. Wegener also knew as a geologist that no single glacier could create moraines on opposite sides of the Atlantic Ocean. And he also knew, as a matter of commonsense knowledge rather than geological knowledge of the possible behaviours of piles of rocks, that moraines do not get up and walk. If bits of one and the same moraine are found in two very different locations, so that one or the other or both of the bits of moraine have moved, then the ground underneath those two bits of moraines has moved. Thus Wegener could offer a reason for thinking that Africa and South America had moved apart. Moreover, when Wegener used one old moraine to determine how Africa and South America had formerly touched, his conclusion matched the conclusion he could reach in a similar fashion by appeal to other, quite separate, old moraines. And the agreement extended to other kinds of evidence that likewise suggested continental drift, for example palaeobotanic, palaeozoological, and palaeoclimatic evidence, and also evidence concerning present-day distributions of plants and animals.

In 1924 Wegener wrote as follows:

> It is just as if we put together the pieces of a torn newspaper by their ragged edges, and then ascertained if the lines of print ran evenly across. If they do, obviously there is no course but to conclude that the pieces were once actually attached in this way. If but a single line rendered a control possible, we should have already shown the great possibility of the correctness of our combination. But if we have *n* rows, then [the smallish] probability [that the match is mere coincidence] is raised to the *n*th power [and thus becomes very small indeed, so that the probability that we infer correctly becomes correspondingly close to 1].
>
> [The match of features between the continents] reminds me of the use of a visiting card torn into two for future recognition.

Notoriously Wegener did not fully succeed with his arguments. He was not able to convince the geological community as a whole to accept the reality of continental drift. An outstanding puzzle concerned how the lighter-weight, physically weaker materials that predominate within the continental landmasses could possibly drift within the far denser, physically harder rock of which the oceanic crust is chiefly comprised. Ultimately this puzzle was resolved when, in the 1960s, geologists began to explore and interpret a new kind of data, specifically palaeomagnetic data concerning the "magnetic anomaly" in rocks on the sea floor. Under a then quite speculative understanding that the polarity of the earth's magnetic field periodically reverses, and that the sea floor freezes the local magnetic orientation at the time when it is volcanically formed, geologists swept the ocean floors with magnetometers. Their object was to discover patterns in the magnetic anomaly in the rocks. Patterns were certainly there, so much so that the regions of like magnetic anomaly proved, on inspection, to be organised in very telling ways. In many regions of the sea floor, bands of positive and negative anomaly were organised, with appreciable clarity, symmetrically on either side of the deep ocean trenches, suggesting that in such a trench new seabed is actually formed and then shunted by new formation behind it laterally outwards in either of two opposing directions at right angles to the ocean trench. Elsewhere, the palaeomagnetic evidence revealed clear indications of subduction of one part of the sea floor under another, or of two facing parts together, or of a part of the sea floor under continental land mass, and so on. Geologists quickly worked out ways to read the patterns for understanding of how the land masses, now conceived to be shunted around by the activity of the sea floor, had moved through ages past. The data, when used in this way, were good out to one hundred million years to the past (and in some regions more). Beyond then, the lines of seafloor magnetic anomaly that could tell a further story would mostly have already been subducted somewhere. The remarkable thing was that the reconstructions steadily outwards to one hundred million years ago agreed in detail with Wegener's findings about past positions of continents.

Through these developments, geology proceeded onto an especially sure path theoretically. In its present activities it is an evidently purposeful and intelligently directed activity. It would be pretty silly now, if it wasn't already silly long ago, to question whether geology is even possible as a science. Thanks to major developments not four decades old, geologists have a very healthy discipline, theoretically speaking, and they know it. What they know is that, by the lights of the ideas now ascendent in their discipline, vast ensembles of seemingly utterly disparate kinds of data are in fact quite excellently well fitted to one another. For example, geologists can chase the continents about into remoter and remoter past historical configurations using inferences from seafloor palaeomagnetic evidence, and find them touching in ways and at times and at latitudes etc. which are required by the palaeoontological, palaeobotanic, palaeoclimatic etc. evidence which originally convinced Wegener and others of continental drift. These are powerfully important developments, and a philosopher's theory of science would be a poor one if it could not do justice to them.

Yet I believe that philosophers prevent themselves from doing justice to such developments, when they begin with the most general questions — how is science at all possible, how can we justify induction, how is a theoretical contention ever warranted by experience — and seek by analytic means at the very outset of their inquiry to answer such questions. To understand how geology as a science is at all possible you need to have your head thoroughly into science. To appreciate for what it is the empirical warrant that there is for the theoretical contention that the continents have moved, you again need to have your head into science. And if you think as a scientist does then you will find no single reason, but rather a symphony of partial supporting reasons, many of them very powerful however, for concluding that the continents have moved. It is the way the many reasons concertedly come together that impresses you. You are impressed not only with this or that analysable reason for thinking that the continents have moved, but more truly by the integrity before reason, i.e. by the synthetic coherence, of the overall ways of thinking and inquiring that bring the many such reasons together. If you have any question to raise about the defensibility of induction, it cannot concern this or that isolated theoretical inference, but rather something concerted. For, your scientific thinking is forever stitching theories together rather than pulling them apart.

§4. *Concerted* induction.

Hume in no way considered the situation of an investigator whose head is already thoroughly into science. That was not Hume's own situation, and in any case he did not consider it. Instead he in effect invited us to examine how we could warrant our very first inductive expectation. He placed importance on analytical focus, and wanted an answer to the most general questions before he was willing to take another step. In that way he came up short and accordingly arrived at a sceptical conclusion.

It is of course only ever in practically idle moments of pure philosophical reflection that we can wonder seriously what reason we have to infer from our past experience any particular expectations at all about the future. As Hume himself points out, the remarkably simple foundational puzzle about warranting induction is without any influence on our practical conduct. However puzzled we may have led ourselves to be in our philosophers' minds, when we stand up from the armchair we do so without the least diminution of our inductive confidence that the floor will support us as it has in the past.

Hume's own explanation for our undiminished confidence appeals to habits and psychology. The explanation that he gives is naturalistic. This firmly sets the matter beyond the purview of reason. Hume singles out past experiences that we have had of the floor supporting us. He contends that an automatic mental mechanism does its work on data like that. In light of the mechanism of the human mind it is only to be expected that we who have in the past been supported by the floor will expect the floor to support us again now. Hume's account is wholly without merit as a contribution to the science of the mind. It is also tail chasing, for just as it invites consideration of our expectations about our expectations, so it invites consideration of our expectations about those and so on. It accounts for our causal beliefs in terms of the causes of those beliefs, which only invites a further such accounting and so on ad infinitum. Yet Hume actually rejoices in the regress. It puts reason in its place and leads us to scepticism, a result which Hume endeavours to show is compatible with common sense. Hume thinks simply that we will have our inductive expectations. We will have our causal beliefs. We are unable to discover good reasons for them, but that in no way blocks our continuing to have them.

My angle on induction is completely different. For I believe that there are multitudinous, richly interconnecting reasons for us to expect that the floor will support us as it has in the past. We are insiders to the situation at hand, and this means that the practical confidence in question is integral to who we are, and that fact connects it to our very mental health. I have *moral confidence* that the floor will support me. Integrity is the watchword for this confidence that I have. My very being would be troubled before I could rescind this expectation. Like any issue about integrity, I cannot put down to a merely particular reason why I am morally confident as I am. And yet my expecting the floor to support me is an issue concerning my rationality for all of that. No analysable mechanism could adequately explain either how the multitudinous interconnecting reasons form or how they can so very effectively crowd in upon me that, unless I seriously doubt myself, I am quite unable to doubt that the floor will support me when I stand on it.

Our confidence (often called *moral* confidence) regarding the support of the floor has a quality that is readily missed under analysis. There is simply no way for us to state *the* reason why we should expect the floor to support us. Reasons in support of this practical expectation are not only too numerous for any one among them to stand out as all-important. They are also inextricably linked with one another, so that the singling out or making explicit in isolation of any one reason cannot but weaken its true purport.

To recognise the synthetic qualities of our moral or practical confidence that the floor will support us is no flight into psychologism. It is, on the contrary, the very reverse. 'Psychologism' is a word used by Kantians, Fregeans, positivists, and Popperians, to rebuke those who confuse issues of fact with issues of right. The synthetic qualities of our moral or practical confidence that the floor will support us connect with genuine issues of right. This moral or practical confidence we have that the floor will support us is rightful, owing to the integrity that there is to it, and thus to ever so many reasons that can be given for it. We can usefully analyse the situation, and by such effort we can begin to see what the right comes to and why it is robust. But because there are synthetic qualities to our moral or practical confidence that the floor will support us, analysis, no matter how thoroughgoing, can never fully discover to us why we are right to possess such confidence. Analysis seeks to *give* the reasons, such as they may be. Indeed, nothing but analysis can truly *give* reasons back to us in the sense of making their character explicit. Yet under any analysis that I can actually deliver, no matter how thoroughgoing I endeavour to make it, my rightful confidence that the floor will support me will be far from completely given. It is impossible so fully to explicate the reasons for a practical confidence that the explication *gives* their full force and character. But that does not mean we are unable to study them as *reasons*. On the contrary, it is often easy to remark why there are reasons upon reasons upon reasons in support of such practical confidence.

Returning, now, to science, I hope that the following is clear from the brief discussion of geology. We need to redress some common deficiencies in philosophers' choices of, and ways with, examples. There is a difficulty about the giving of examples that it is important to understand. We are explicit in how we give an example. This severely limits what there is about it that we can then place under analysis. In my view this problem is acute both in moral philosophy and philosophy of science. Moral discernment is a capacity that far outstrips our wherewithal with mere words or thus our ways of being explicit about our reasoning. So when a philosopher sets before us an example for us to think about morally, our capacity for moral discernment will either send us beyond the explicit description to merely possible nuances that the scenario could have had, or it will cleave to the explicit description and become something severely stunted in itself. Either way the example is under its explicit description not liable to help us appreciably to understand the qualities of our own moral discernment. For science, the situation is similar. A philosopher may invite us to consider an induction of the form "swan one is white, swan two is white, … (all the way up to), swan fifty-seven is white, therefore all swans are white". It is supposed to be a virtue of the example that it is in its every salient characteristic completely set before us. We can therefore go to work on it analytically, and assess whether the inference is rational. If it is in the least way defensible then we will be able to identify, indeed give ourselves, under the analysis, just what the defensibility is. Otherwise we will conclude that it is indefensible.

In fact these expectations are naïve, as anyone with the least scientific discernment will readily see. For someone using science would immediately add collateral considerations and discern potential richness to the inference. What we might infer about a colour of all swans from an experienced sample in which all were white we would infer on the basis of antecedent understandings that we have of the aetiology of various characteristics that animals might have and thus of what might *cause* the characteristic in question (whiteness) to be endemic to swans if it is. Given the way inheritance works, and the common heritability of surface colour, and the known uniformity of experienced swans so far, it is, we might judge, possible, but hardly certain, that all swans are white. Knowing what we know and seeing the uniformity in the sample so far we feel a palpable urge to generalise. We are however easily able to discern why the sample could be as it is without the generalisation being true. So if we generalise we will do so tentatively, with little confidence. By contrast, if we observed that swan one had a heart, we would not need to look any further than that to infer that all swans have hearts. Indeed, if swan one were observed to bleed, we could with almost equal safety infer that all swans have hearts. We discern an impossibility here, from knowing what we do, that any swan could be blooded without them all being blooded, and that any blooded creature could lack a heart. Moreover, in quite the other direction from the "all swans are white" inference, we could consider the case where swan one has a wart on its left eye, swan two has a wart on its left eye, … (all the way up to), swan fifty-seven has a wart on its left eye, and thus all observed swans have a wart on their left eye. We know enough about the aetiology of warts to know how foolish it would be to infer from this that all swans have a wart on their left eye.

The three swan examples are as given formally the same, but there is a world of difference between them. So, so much the worse for hoping to bring all the salient considerations into view by explicit description of an example.

I have said that only ever in practically idle moments of pure philosophical reflection can we wonder seriously what reason we have to infer from our past experience any particular expectations at all about the future. In any actual real-life practical situation we invariably have a veritable concert of reasons for making such an inferential step one way rather than another. Of course all those background reasons are themselves inductively grounded. But the point remains: it is never the one enumerative induction we perform, but always something concerted. Induction is never enumerative, it is always concerted. And as I have begun to make out, there are synthetic qualities to concerted induction that make a world of difference so far as the rational justifiability of such inductions is concerned. Under analysis, concerted induction proves to be rich with deductive features. But the overly celebrated idea of "hypothetico-deduction" quite prevents our discerning these features. To appreciate as far as we can in terms of deductive features how good inductive inferences are made, we need the idea of a measurement inference. I have illustrated this idea by discussing the way that geologists adduce empirical reasons for thinking that the continents have moved.

It does egregious violence in many ways to the actual activity of science to fashion 'all swans are white' a scientific theory. Two groups of philosophers were once very prone to illustrate scientific theorising by considering such examples. These were the logical positivists, and the Popperians. The logical positivists adopted a verificationist outlook, according to which a theory in science, insofar as it is meaningful and true, *accurately predicts what we observe*. Popper preferred to leave aside the question of meaningfulness and concentrate instead on that of demarcating science. In effect he preferred the formulation that insofar as a theory is *scientific* and true, it accurately predicts what we observe. Contemporary philosophy of science corrects these views at least to the extent of replacing 'accurately predicts' by 'concertedly harmonises' and 'what we observe' by 'phenomena'. A theory of science insofar as it is meaningful and true *concertedly harmonises phenomena*. Or at least, a theory insofar as it is scientific and true *concertedly harmonises phenomena*. A phenomenon is as far different from what we might observe as concerted harmonisation is from mere accuracy of prediction, and the differences are very similar. Phenomena have a richness far and away beyond what can be brought under simple observation, yet phenomena also possess robust consilience features which remark a kind of harmony in what they draw together. A good example of a phenomenon is that the floor will support me when I walk on it, and that is precisely the sort of explanandum that scientific theorising is meant to harmonise concertedly with other phenomena and thereby explain.

Most phenomena for which science seeks systematic explanation are delineated by experimental means. Measurement, the practical activity of deducing a theoretical conclusion from empirical data, typically enters into the experimental work through which a phenomenon is delineated. Indeed in the delineation of any single phenomenon measurement typically enters not once only but rather many times over, in a rich, textured, multi-level sort of way. Because this is so, the very distinction between data and phenomena is itself layered and contextual, with every phenomenon both serving to colligate (and thereby qualify or interpret) an assemblage of underlying data, and in its turn itself standing as datum for one or more higher-level such colligations. You have to look a long way down through the layers of phenomena colligating data which in turn are phenomena colligating data and so on, before one encounters data that at any stretch can be called directly observable. And even at that extreme it is naive to call them directly observable without qualification.

An amply measured experimental finding is rich with consilience features because any one aspect of the finding can be measured in a large number of different ways that appeal to quite different theoretical assumptions and quite different ranges of data.

Yet that same richness of engagement of evidence is to be found at the higher level of theory. Theories *harmonise* the phenomena they explain, they do not merely *predict* them. Aspects of the theory can be deduced from phenomena (i.e. a theoretical explanans can be deduced from one of its own explananda, in a way that uses background theoretical assumptions of course), and not in one way only but in many ways. That is to say, background theoretical assumptions, often themselves quite weak and innocuous, can be used to show the salience for further theory of a certain experimentally adduced phenomenon. But the same lesson for further theory can be learned in other ways as well, using quite other background theoretical assumptions, and quite other phenomena. This rich engagement of phenomena by theory amounts to consilience. Because of it, overall theory reveals a kind of hidden harmony in the diverse phenomena. That is to say, overall theory richly harmonises the phenomena. The demand that a theory should have this kind of relation to evidence is vastly stronger than the demand that a theory should merely accurately predict the phenomena. And the difference closely resembles that which there is between phenomena themselves and the paler empirical issue of "what we observe".

All this to say that ‘all swans are white’ is an egregiously poor illustration of a scientific theory. It is worse than a toy example: it completely misleads, by its lacking all the synthetic qualities that a scientific theory worthy of the name invariably has. It also helps fate philosophers radically to underestimate the logical richness of the thinking that goes on in science.

Such synthetic qualities to our thinking in science can to a considerable degree be illuminated. Calling them ‘synthetic’ is appropriate but this draws no veil of mystery over them. One kind of analytical approach that is helpful is to seek to examine the myriad ways in which a well considered theory is measured. To say that the theory has synthetic qualities does not imply that the analysis of the reasons in support of the theory has limits beyond which it cannot pass. It implies, on the contrary, the very reverse. We can analyse for as long as we like, and still there will be further reasons left to uncover. Our only mistake would be to expect that by a single analysis we can reveal in its entirety the reasons in support of the theory. That mistake would egregiously underestimate the worth of the theory or the rightfulness of our conviction in it.

§5. The analytic approach itself the source of the philosophical confusion concerning induction.

Induction has been called "the glory of science and the scandal of philosophy". It will be clear by now that I agree with this assessment. And I blame the long history of philosophers' attachment to, but difficulties in answering satisfactorily, the ostensible general "problem of induction", upon their expecting the answer to be a point, upon which an analysis can be focussed. Thus I believe that:

(8) **Philosophers' problems about induction are largely due to their own unremitting disposition to be analytical.**

In any actual real-life practical situation, including, most assuredly, any actual juncture in the development of a science, we invariably have a veritable concert of reasons for making our inferential steps one way rather than another.[2] Thus in order to feel exercised in relation to science by the supposed problem of induction we would have to bracket all the discernment of which a scientist is

[2] The rare exceptions prove the rule. If in a given situation our discernment seems to fails us so that our situation seems to us ambiguous and we are not drawn by our experience so far to form any expectations at all, this may be because two roughly equal concerts of reasons support our making opposite inferences about the future. Alternatively it is because, try as we might, we cannot find a foothold anywhere within our present system of beliefs for judging the situation at hand and projecting forward from its features in any reasonable way. Either way we are reminded that usually we do have a synthetic foothold and a myriad reasons for projecting one way rather than another.

capable. Yet such discernment or concerted reasoning (generally inexplicit intuitive reasoning, that if we stopped long to articulate and analyse would prove to be highly ramified in what it involved) will invariably be special to a given practical and epistemic context. And significantly, in order to *give*, in the sense of perspicuously analysing, any one of the abundant reasons for making an inferential step one way rather than another, we must make that reason out as after all deductive in form. Under analysis, such reasons will reveal themselves as valid deductions to a theoretical conclusion, from premises that involve not only empirical elements but also a background of other theoretical beliefs. Under analysis, such reasons will, in short, be from measurement inferences that we perform. What is reasonable about the induction in question will come down to how carefully measured it is. Of course all those background assumptions, without which it is impossible to measure anything, are themselves ultimately only empirically grounded. So in some sense we thus see induction, rather than deduction, ever more in the picture that we need concerning our epistemic behaviour, rather than less. The synthetic activity of bringing together many such measurement inferences and then on the basis of the whole concert of them inferring the theoretical conclusion to which they point, is also notably non-deductive. In fact its many deductive qualities are so formed not explicitly but typically only under our analytical reconstruction.[3] Our inference making is almost all intuitive, so that a great analytical reconstructive effort is required to formalise it well, or thus make it out as strictly logical. Without any doubt therefore any *measured* view about anything is in part a *synthetic* achievement.

The point remains, however, that it is never a single enumerative induction we perform, but rather, in our deductive reliance on a background of already theoretical beliefs, always in effect a *concerted* induction. Thus

(9) **Induction is never enumerative, it is always concerted; but, whatever is concerted can never be fully brought into view under analysis.**

Why, in reasoning inductively, we do all right, admits of no single simple analysis; and yet, in the specific instances of our reasoning thus, the integrity of what we do will often nonetheless be robust, that is to say, both rich and thoroughgoing. The more we analyse our reasons the more we bring this to light; but, we can never

[3] Similarly but more specifically, when we formalise a deductive argument in order to analyse it, we reveal grammatical features that were only tacit, that is to say, not at alll part of the surface grammar of the argument as it was presented in everyday language.

complete the analysis, nor thus ever fully articulate all the reasons that we have.

It is true that some philosophers of science, following Popper, seek an actually anti-inductivist understanding of how there can be science. That is, they seek to show that there can be science not by defending induction but rather by pointing to what they suppose could be a wholly anti-inductive kind of intellectual activity. They suppose (falsely in my view) that scientists should be able, in an intellectually rich, systematic way, to entertain and employ theoretical contentions quite without supposing that any of these contentions is ever in the least way actually warranted. Popperians call this position 'critical rationalism'. I believe that critical rationalism fails as a philosophy, and that the disposition to uphold it proceeds in any case from a philosophical mistake. That critical rationalism fails as a philosophy has, I believe, been adequately argued elsewhere.[4] To the extent that I concern myself here with critical rationalism, I intend simply to explain my reasons for saying that the disposition to uphold it proceeds from a mistake. The mistake in question infects other philosophies of science as well; it is not the special burden of critical rationalism. Instead it is, for reasons

[4] See Putnam, Salmon.

that I have explained, a mistake that follows all too readily simply upon the adoption of the analytical orientation in philosophy. I believe, with Popper, that analytic philosophy on its own cannot honestly and searchingly reckon with induction without ultimately reckoning induction ineluctably problematic and unwarrantable. But this, I believe, is the very signature of a mistake. If induction seems ineluctably problematic and unwarrantable, then this epitomises the way that philosophy can grow silly when it is freighted with an unremittingly analytical orientation. I want to be careful however not to overstate the extent of my misgivings about analytic philosophy. Its methods are in fact crucially important for illuminating how inductions are warrantable and thus how science is possible; their drawback I shall argue is only that they can never by themselves complete these tasks. Whenever a theoretical inference is warranted, it is warranted for *countless* reasons, any one of which we can identify however only to the extent that we can adequately analyse it, and that only to the extent that we can adequately make it out under careful analysis as after all deductive in form.

In connection with this image of inductive reasoning as concerted, I choose to call even Bayesian probabilistic reasoning “deductive”, inasmuch as the output probabilities can be *calculated* given the experiential input and the priors. Some philosophers who intensively study probabilistic reasoning choose to label such reasoning 'inductive' in contrast with non-probabilistic deductive reasoning. I believe that this common kind of appropriation of the word 'inductive' is a mistake, every bit as much as is the aspiration tightly attached thereto, to proclaim some suitable theory of probabilistic reasoning to be both a successful analysis of induction in general as well as an adequate general explanation for its warrant. In fact the analytically styled work by probability theorists becomes a helpful but relatively special part of the overall understanding we need of empirical reasoning, and in my view it is best to consider the probabilistic inferences thus analysed as actually but a part of the overall nexus of measurement-inference *de*ductions. Probabilistic reasoning using empirical evidence can be one way in which we render our theoretical ways of thinking carefully *measured*. Bayesian probabilistic reasoning resembles non-probabilistic measurement inferences in the way that it depends on engagement of a wider nexus of already theoretical beliefs, and in the way that it is, when analysed, formally licensed by a calculus. Bayesian probability reassignments on the basis of evidence are therefore in their way acts of measurement. Thus to the extent that the warrant for a theoretical contention has been analysed by Bayesian considerations, so that the theoretical reasons have been laid out for our accounting both certain prior and certain conditional probabilities a certain way, and thence for drawing from an empirical consideration a certain revised level of confidence in the theoretical contention in question, I say that the warrant that has been thus illuminated for this theoretical contention is deductive in form. And I mean to include such probabilistic reasoning when I insist that the reasoning that we can analyse and thus make out as deductive in form is only a proper part of that which we concertedly perform when we induce in a considered way a theoretical conclusion. Why we cannot dispense with induction in favour of a deductivist reckoning of science is in my view simply that analysis can never completely illuminate the warrant for the theoretical inferences that we make. It may bring into view, more or less clearly, a myriad more or less cogent partly empirical reasons in support of the inferred theory. But the scientific warrant for the theory in question will invariably come to more than can thus be brought into view. I say this even in light of my broadening of the scope of the term 'deductive', to include reasonings that conform with an explicit, rational canon of probabilistic reasoning rather than a logical canon of reasoning that is non-probabilistic and thus fully demonstrative.

§6. Induction from the standpoint of analytic philosophy.

The philosophical study of the analytic function of reason is logic. Around the beginning of the twentieth century, logic in this sense profoundly changed. Without a doubt, the advent of a new logic was a great step forward for philosophy. It is pertinent to pause briefly over this development. This development helped establish some enormously high levels of expectation upon subsequent philosophy. The standard that it set was in some ways a benefit to philosophy. But the high expectations were also in some ways artificial and naïve. By examining this history we can learn something about why induction, in contrast with deduction, registered during the twentieth century as a deep and insoluble problem for philosophy. We can learn why Hume became especially influential.

The following comparison is useful. The development of modern logic was one breakthrough. Another (later) breakthrough in human knowledge that I wish to compare with it was the discovery of a molecular basis for genetics. Each of these breakthroughs established a high level of expectation. That high level of expectation was not entirely misplaced, but it did greatly over-estimate the pace and extent of the further accomplishment that there was liable to be. Each of the two breakthroughs in its way clearly told us that already extant sciences were in great shape. People had had to become wonderfully well sorted out in their thinking about mathematical analysis before anyone could mine that well sorted out thinking for a new and improved understanding of the nature of deduction itself. Likewise people had had to become wonderfully well sorted out in their thinking about various facets of life, inheritance, chemistry, and physics, before there could be a discovery such as that of a chemical basis for inheritance leading up to the breakthrough by Francis Crick and James Watson. As soon as these great breakthroughs were achieved, however, everyone's attention was less on what went before than on what might come after. People seemed less concerned that science had had to be wonderfully well sorted out already in some ways in order for the breakthroughs to have happened at all, than they were that these breakthroughs might begin a new productive phase of further investigations. In their excitement about the breakthroughs people formed expectations about the future that in each case later proved impossible to fulfil.

With the advent of modern logic, philosophy seemed at first blush to have broken through to a powerful, topic-neutral, totally general understanding of the nature of valid deductions. With the advent of a theory of DNA, life science seemed at first blush to have broken through to a powerful, topic-neutral, totally general understanding of the nature of genes. Those were the expectations, at any rate, and initially people had some good reason to think that even such new and high expectations were now well on the way to being fulfilled. The answer about valid deductions was framed in terms of a canon of reasoning, and that about genes as a theory of the structure and functioning of molecular DNA. Just as any valid deduction from whatever sphere of thinking was expected to fall appropriately under the anointed canon, so any gene, from whatever living species, was expected to be a stretch of DNA. The canon of reasoning was expected to illuminate all science. The DNA code was expected to illuminate all life.

Of course the new logical canon was employed more or less straightaway by workers who successfully went on to illuminate some specific things about scientific reasoning, just as the new reckonings in molecular genetics were used more or less straightaway by workers who successfully went on to illuminate a few specific biological phenomena. The successful work of this sort in philosophy of science displayed the scope of deductivism. The successful work of this sort in life science showed that some identifiable biological phenomena chiefly express the presence and biochemical action of DNA genes. But working up new knowledge in these ways proved to be a whole lot more effortful and limited in scope than initially expected. And the successes such as they were sat alongside other developments that went still less according to plan. Philosophers began to proliferate variant logical theories and in any case to downsize deductivist expectations in the philosophy of science. Life scientists adduced reasons for denying that all genes are DNA, and in any case they downsized the level of expectation surrounding genes in biological explanation. Logic, though important, could not be held up as the very key to understanding science. DNA, though important, could not be held up as the very key to understanding life.

It is a gross exaggeration, but remains even now, because of the beauty of the theory, an almost forgiveable exaggeration, to suggest that, say, first-order logic with identity provides a complete topic-neutral understanding of the difference between a good deduction and a bad deduction. The very high level of expectation that once surrounded logic engendered the expectation that a theory of *in*duction would, if acceptable, likewise canvas once and for all in a complete, topic-neutral way what the difference is between a good induction and a bad induction. Unless that level and quality of understanding could be mustered (as it most certainly cannot), philosophers felt that they had a problem about induction.

In fact deductive logic (however broadly construed) provides but a narrow window onto human intelligence. Logic on its own is quite unable to tell us how it is that reason is meet either for organising our practical wherewithal or for interpreting experience. Logic on its own cannot completely tell us how we generate scientific knowledge. That limitation upon deductive logic links to why we need to countenance inductive reasoning after all. And significantly, philosophy of science is no longer hampered as it once was by the narrowly limited scope of logic. Philosophers of science have helped broaden beyond logic our understanding of rationality. They have paid ever more attention to practical dimensions of the rationality of science, and thus to questions, not altogether within the purview of logic, of the synthetic integrity of ways of being or acting. They have helped discover how even a way of *thinking* in science depends for its coherence or rational integrity and also its content or meaning upon a gathering together or harmonising of *practice*. Deductive logic plays a key role in the harmonisation of theory and practice, but a role that would be impossible if our rationality were not at the same time broader than logic in its purport.

During the headier days for logic and analysis in philosophy, many philosophers declared that Hume had in many ways been right. Hume worked in the wake of Newton and thus when science had itself at last been established as a going concern. And yet, Hume made it seem actually questionable whether science is at all possible. Of course, this makes it philosophically very interesting and worthwhile to consider Hume. The questions whether and if so how science is at all possible are philosophically very fetching. They are, in philosophy of science, the most general questions of all, and for this reason they are especially apt to captivate the philosophical imagination. Philosophers of science for some while attempted to adopt Hume's analytical orientation, thus agreeing with Hume on many points. They endeavoured nonetheless, pace Hume, to explain science as possible after all. But in all their early attempts they were not successful. Then gradually, over many decades, they were moved by their work ever further away from the very starting points of Hume. They have in many ways discovered the worth of a synthetic orientation in philosophy, and learned only as they have thus departed from Hume in what way to explain, pace Hume, how science is even possible.

From the whole development we may learn the following, very important, meta-philosophical lesson. Philosophy should not necessarily *begin* by answering the most fetching questions. Often it needs simply to be philosophy's *end* in the long run to answer these most fetching, most general questions of all. To attempt to answer how science is at all possible, we could insist upon analytical focus. We could tell ourselves not to take a single further step until that question is answered. We could demand our answer to be simple and sharp. But this would all be a mistake. We would never get anywhere that way. Instead, we should set out in a number of directions, and expect in the end to find ways to put many considerations together. It is, by default, the latter path that philosophy of science ultimately took. Only after it did so was it able to answer how science is even possible.

§7. Induction from the standpoint of synthetic philosophy.

The analytic orientation that produces for Hume the problem of induction also leads him to treat the kind of cognition we marshal in doing science as one thing, and that from which we derive *moral* or *practical* direction as quite another. Thus Hume's moral philosophy concerns what is supposed to be a different and significantly separate facet of our cognition from what we employ when we do science. I, however, trust neither the moral philosophy of a thinker who, like Hume, concludes that empirical science is not truly possible, nor the philosophy of science of a thinker who, like Hume, is indifferent to our rational freedom. What relates science on the one hand with ethics on the other is the partly synthetic function that reason has. Hume concerns himself with only the analytic function of reason, placing the synthetic capabilities of the mind, which he grossly oversimplifies, in the unreasoned realm of natural, "associationist", psychology. Yet when we examine the practical connections of actual scientific theorising we are got to see quite sharply the partly synthetic function that reason has. Indeed we see the way that, both in science and in the larger practical sphere, the logical or analytic function of reason itself works significantly in support of this synthetic function.

By their every well reasoned inductive contribution, scientists rationally synthesise new connections between theory and practice. In ways that can be more or less dramatic, but in typical work are too subtle to remark, scientists thereby rationally alter the concepts with which they were antecedently possessed, rather than merely following them. For, the conceptual fabric is determined to be what it is precisely by the way that theory synthesises practice. Creative geniuses such as Albert Einstein can so affect this fabric that they significantly change the understanding of a whole community of scientists about the ostensibly elementary facts. Yet the changes they wreak are rationally required ones in the light of varied evidence to which they can point. The better we understand an Einstein's way of 'scenting out the path that leads to fundamentals' as he put it, and of adducing whole new ways to think in a science, the better such work illustrates for us the following fact. *The highest spontaneity of the understanding is also the most responsible, and so produces concepts that are already richly supported by evidence.* The imagination is free, but the deliverances of the *freest* imagination are in fact the *least* conjectural.

In the wider practical sphere, people are likewise the more free the more responsible or measured are their judgements. For only to the extent that one acts in a reflective, rational, and thus measured way, is one's determination really from and of oneself. The *synthetic* qualities of true rational self-determination are palpable. Rationally reflective practical self-determination *reshapes* a person's desires. It does not merely sort these desires logically and then follow them. In most cases, the reshaping is too subtle to be remarked; and yet, whenever people rationally resolve a big ethical issue in a truly creative way, the fabric of their desires is reshaped quite dramatically. And this kind of rational self-authorship and reworking of a situation evinces the same positive quality of freedom that I remarked above. It is maximally responsible, and thus maximally determined or necessitated rather than merely negatively free, that is, merely unconstrained. Creative geniuses such as Ghandi can affect the understanding of a whole society about what the elements are of everyday situations and about which among these elements have worth and why. Yet the wonder about these creative acts is precisely the outright necessity that we are able, looking back, to discern in them.

Ethics and science are orthogonal to one another it is true, the one concerning the rationally best or most harmonious and elegant way to consider our situation as we empirically or practically find it, the other concerning the rationally best or most harmonious and beautiful way to fashion our situation to the extent that we are practically able to create it or actively condition it. But in their orthogonal conformation to one another science and ethics are not wholly disconnected, for each is regulatively conditioned by synthetic rationality. Hume's analytic philosophy unhelpfully presents neither science nor ethics as truly rational. Just as the scientist is reduced on Hume's analysis to following rationally unwarrantable animal habits, so even those actors we would call morally most sensitive are in Hume's view merely products of the past conditioning of their desires or passions. No true rational creativity, no positive freedom, can be remarked in their conduct, for they only follow their own extant sympathies, sympathies that either are merely inborn to them or have been worked up in them by the social processes that have impinged upon them. In fact all moral discourse is itself just a part of the fabric of these efficient causes. Hume offers to explain to us why it actually has no objective content, and why we need an error theory to explain our own ineluctable disposition to think that it does.

The Hume believes that both in the way we develop science and in the wider sphere of practical action we are all just the sum of our personal past experience. This view quite fails to reckon with the agency that there is in experience itself. We cannot be just the sum of our personal past experience, because what our experience itself amounts to depends in part on how we ourselves synthesise it, under concepts of our own making. Of course if we were just the sum of our personal past experience, then our synthetic or rational freedom would be illusory. That is Hume's belief, but this belief is less the conclusion of a convincing argument than it is built into Hume's very starting point.

These two points of Humean irrationalism, that about science and that about morals, follow from his analytical orientation, by which he wholly removes the synthetic function of reason from view. They are the very signature of a mistake.

Thus I hold that

(10) **Just as reasoning has *synthetic* aspects of which Hume takes no adequate account, so philosophy is *syncretic* to an extent that Hume quite fails to register. In particular, although science and ethics are orthogonal to one another, they are in that conformation connected. We cannot hope to attain adequate philosophical insight into the possibility in general of theoretical knowledge such as we find in science, without achieving similarly adequate philosophical insight within practical, or moral, philosophy; and conversely. In both**

science and ethics there is a synthetic role for reason which Hume quite leaves out of account.

Philosophy has long polarised its thinking about thinking around what in fact are complementary understandings of reason itself, the one understanding analytic, explicative, and symbolically or logically oriented, the other synthetic, dialectical, and practically or ethically oriented. Prior to the Scientific Revolution this polarity largely agreed with that between empiricists and rationalists. Science came into its own however precisely by compounding and accentuating the practical or experimental connections of theorising in ways that surmounted the very divisions above-mentioned. Unfortunately the empiricist and rationalist camps that continued in philosophy, although they both accepted that the Scientific Revolution surely represented some kind of advance, each did not sufficiently alter its own terms of reference for seeking to understand that advance. One needs to step above the very opposition properly to see how science is possible. The philosophers who have done this are not 'ism' philosophers. They stand above the oppositions between competing 'isms'.

It is not too gross a generalisation to say that empiricist philosophers have often either not known science well or have been befuddled by it, whereas in the early years of science many rationalist philosophers knew it and practised it well. G. W. Leibniz for example was a very significant scientist, not that this leads me to endorse his philosophy of science, for I do not. Kant in the eighteenth century and William Whewell in the nineteenth are further examples of philosophers who had their heads well and truly into science. They are sometimes called rationalists; but, in ways of which there is an inchoate expression even in Leibniz, they actually rather stand above that polarisation, so that we can with equal truth and equal falsity likewise make them out as empiricists. By contrast, I can think of no significant strongly empiricist philosopher prior to the twentieth century who was not scientifically inept. Many of the twentieth century empiricist philosophers were relatively more sophisticated scientifically, yet were dealing with science that had itself grown especially difficult, and so were often actually befuddled by it. In general, with only rare and relatively recent exceptions, empiricist philosophers have not personally known well what it is to have one's head well and truly into a science.

In fact successful science makes rich with practical connection its every theoretical pronouncement. Having one's head well and truly into such a science is, at the level of practical engagement and wherewithal, something indescribably rich. This fact remarks both the truth and the falsity of both rationalism and empiricism. Successful science, as it develops, adduces new theoretical concepts in ways that are *both* empirically measured *and* notably dialectical. By 'dialectical', I mean that these ways of adducing theoretical concepts disturb the pre-existing fabric of concepts, and thus disturb even how people henceforth interpret their experience. Yet, pace Thomas Kuhn and his followers, there are no real revolutions in science. Scientific progress is palpable; successor conceptions are superior in their engagement of evidence to their predecessors. If rationalism prioritises organising conceptions to particular facts, and empiricism prioritises particular facts to organising conceptions, then by pursuing the most measured understanding of things science transcends the very opposition between rationalism and empiricism. For in an experimental measurement organising conceptions are prioritised to particular facts in the sense that scientists employ their antecedent theoretical considerations to identify what particular facts would be telling and in what ways. Yet measurement also prioritises particular facts to organising conceptions, in that the measurement outcome is telling precisely about some further reach of theory. Because of this duality, the quest for ever-more-measured understandings in science is in some degree hermeneutical—openly interpretative—and thus to the same degree notably dialectical. Thus not only rationalism is telling yet false, but also empiricism is telling yet false.

The very opposition just discussed between rationalists and empiricists also concerns the synthetic versus the analytic style of philosophical investigation. Empiricism promotes an analytic style of philosophical investigation; rationalism promotes a synthetic style. Philosophers of science have, in effect, during recent decades, both refined the analytical orientation and also discovered reasons significantly to relax this orientation. Indeed I would say that philosophy of science both led philosophy across much of the English speaking world into a profoundly analytic phase, which peaked in its singularness around the middle of the twentieth century, and then more recently has significantly led those same schools of philosophy somewhat back out of that purely analytic phase. Two illustrations of the latter development are the renewed interest and respect for Kant as a philosopher of science as well as Whewell, and all the work through which contemporary philosophy of science has improved its understanding of the practical connections of scientific theorising. In measured but significant respects in their investigation of meaning, mathematics, theory change, and experimental science, contemporary philosophers of science have employed ever more the synthetic, not only the analytic, stance in philosophy. It is this that makes it possible at last to overcome once and for all the traditional conundrums concerning induction that Hume set in place.

In a way that simply wears on its sleeve the synthetic orientation that I say has staged a resurgence in philosophy, I take the philosophical route past Hume's conundrums about induction to be almost as significantly about moral philosophy as it is about philosophy of science. Thus I follow the rationalists somewhat far into their syncretism about philosophy. Yet we will scarcely discern from such an approach how genuinely to move past Hume's conundrums about induction unless we cut through or avoid the infamous obscurity of the rationalists. I believe that I do this by discussing the nature of measurement. For I believe that

(11) **Both in science and in the wider practical sphere responsible actors seek the most measured understanding possible.**

Yet from this simple idea I believe some important contentions follow, contentions that rationalist philosophers often advance, albeit often with undue obscurity.

(12) **The quest for a most measured understanding is inevitably dialectical. It sets theory into so rich a connection with practice that this connection is never fully analysable. The importance of measurement renders *immeasurable* the full practical connection or significance of any reach of theory.**

Already in the brief discussion of geology I believe I have illustrated these contentions and displayed something of their worth. I will next further elucidate and illustrate these contentions with a supplementary discussion concerning physics.

Consider the assertion: 'metal fatigue diminishes electrical conductivity'. There is no hopeless holism about how this assertion connects to evidence. But the issue of how it can be empirically verified to the extent that it can is bewilderingly rich. We can bring logical clarity to *aspects* of why we are warranted to believe that metal fatigue diminishes electrical conductivity, and every way in which we can do so helps illuminate the warrant for what it is for this belief. But the warrant can never wholly be brought into view, or thus made clear in its entirety by analysis. It is too rich, and overconnected (in multiple independent ways) with too many different aspects of the ways we think and act in the world.

Thus, consider how I might measure the conductivity of a particular metal alloy, both before and after the metal is mechanically fatigued. Notably, "direct observation" is not going to help me. One way for me to measure the conductivity of the metal alloy might be for me to apply a voltage across a sample of the alloy, a sample that has been suitably drawn into a wire of known length and known constant cross section. I then measure both the voltage and the current across this sample wire. For the measurement of the voltage, I use a voltmeter, for the measurement of the current, an ammeter. I must have a well grounded confidence that these instruments are up and running. I use their readings not directly, but rather via a calculation that also takes into account the length and known constant cross section of the sample, to determine the conductivity of the sample.

On what kind of grounds will I base my confidence that my instruments are up and running, that is to say, accurate in the information they give me? Consider the voltmeter. Its design of course reflects theoretical understandings of one sort or another, but it is not specifically by virtue of our confidence, such as it is, in these theoretical understandings, that we adjudge the voltmeter to be up and running, i.e. reliable for the measurement of voltages. On the contrary, the reasons why we judge it this way is richly practical. For a start, it is, in its behaviour, notably alike to other voltmeters, including to voltmeters that are designed quite differently, some so differently as to reflect a largely separate reach of contemporary theoretical understandings. The consilience of the many voltmeters, i.e. their robust agreement with one another, is a pretty fine indication that there is a genuine phenomenon, called voltage, which it is appropriate for us to measure, and take a theoretical interest in. Secondly, however, there are other very immediate practical checks that confirm that the voltmeter I have in front of me accurately measures voltages. For example, I can test it across a "D" cell, and discover that the needle jumps to "1.5" on the dial; or across two "D" cells in series, and discover that the needle jumps to "3" on the dial, etc. A host of practical considerations like these can attest not only to the reliability of the instrument but also to the reality of the phenomenon it helps us measure. A similar story can be told concerning the ammeter.

By combining the measurements of the voltage I have applied across the sample of the metal alloy, and the current, and the cross sectional area and the length, it is straightforward for me to calculate the conductivity. That is how I measure the conductivity. It helps that there are robust quantitative phenomena such as voltage and current for me to lay hold of using instruments. It helps that there are robust regularities such as Ohm's law, for I calculate in terms of these as an important step in my measurement. If there weren't already many measurable phenomena and many mathematically fairly precise relationships between these phenomena, I would not be able to make a meaningful measurement in the way that I do. But what I do is not any kind of step from direct observation.

Suppose that I now mechanically bend back and forth, repeatedly, the very same wire whose conductivity I have already checked. I fatigue the metal, almost but not quite to the point of its breaking. And then I measure its electrical conductivity all over again. The results of all this work would be one empirical instance that is relevant to the question whether metal fatigue diminishes electrical conductivity. I might never have produced an instance in quite the way that I produced this one. I might have used quite different measuring devices, that operate according to quite different principles; I might have selected many samples of metal alloys but not the particular alloy in question. The empirical instance that I have in fact come up with is enormously rich in its characteristics but so also would innumerable other empirical instances have been that I might have considered instead. And just as these potential instances are vastly different among themselves, so also is no one of them strictly needed, in order for me to have warranted, perhaps superbly well, my contention that metal fatigue diminishes electrical conductivity.

When the characteristics of a particular example, such as this one, are sketched with reasonable care, it begins to seem that meaning is empirical and epistemic but in altogether too rich a way to be brought into view all at once under an analysis. There is no difficulty about 'inscrutability' of meaning: the problem if anything is just that there is far too much to scrutinise. It was W. V. Quine who set philosophers to worrying about an ostensible 'inscrutability' of meaning, a worry which he linked to the ostensible problem of underdetermination, which is a direct expression in other terms of the ostensible problem of induction. Quine is correct to suggest that we cannot scrutinise the whole meaning of a term if that would involve our giving in toto *an analysis* concerning the empirical connections for the term of the whole meaning of it. But there are two words here, 'an', and 'analysis'. It is important to remark that the problem word is 'an' rather than 'analysis'. In fact we can give a wealth of different analyses of how a theoretically meaningful expression is empirically meaningful. Each of these analyses is pertinent and illuminating but incomplete. The incompleteness of every such analysis means that there is not an analysis in light of which to say that the meaning has been scrutinised. But the activity of giving many such analyses is illuminating, and I would say that it is an excellent way to scrutinise the meaning of the term. There is much to analyse, so much indeed that we cannot begin to bring it all into view. This is not a problem of inscrutability so much as it is an embarrassment of riches.

[operationalism]

§8. Conclusion: a non-Humean picture of deduction, induction, analysis, synthesis.

Two things should be clear by now: my picture of induction is totally different picture from Hume's, and it is a good deal more apt than Hume's to the situation in geology as I have described that. By employing my picture, we are made much better able to comprehend the integrity of geology as a science. I have briefly examined reasons for expecting that we need my picture in order to appreciate the integrity not only of science generally, but also of much of our everyday thinking about the world.

Inductive inferences are often contrasted with deductive ones. We have seen the need to be careful about the contrast. Inductive inferences have theoretical conclusions, but it is likewise possible for a deductive inference to have a theoretical conclusion. Indeed it is possible for a deductive inference both to have a theoretical conclusion and non-trivially to employ specific empirical propositions among its premises. Inductive inferences are a challenge to analyse but many deductive inferences are too _ for, very often, vast analytical work is required to illuminate the logic of deductive inferences made in a natural language and everyday circumstances. What then is the best way to contrast inductive inferences with deductive ones? I have offered the view that in general an inductive inference is not an argument, and neither is it best modelled as one. So there is no form at all that we can legitimately read into it. By contrast, deductive inferences are best modelled as arguments, possessed of a form.

In a paradigm deductive inference, thoughts that are taken to be a sufficient guarantee that the conclusion is true are explicitly laid out as premises. There are cases, actually rather rare outside of logic classes but signal nonetheless, in which the sufficiency of the premises for the conclusion not only obtains but can readily be illuminated by an explicit, formal, very general canon of valid deductive reasoning, that is to say, by a system of symbolic logic. Around such cases there is a wide penumbra of inferences that also deserve to be called deductive, including enthymemes, and deductions that would need significant work in order to be formalised or thus illuminated by symbolic logic. The qualities of such penumbral deductions, valid or invalid, are extremely various, and it is in my view doubtful whether a single clear overarching characterisation can be provided which explains why they all deserve to be called deductions. (The problem is parallelled and indeed compounded by the evident need to countenance not one single logic of deduction, but rather, in order to do justice to the great variety of demonstrative inferences, many such logics.)

All remaining inferences are, since they are not deductive, inductive. Consider the entirety of this collection. It again does not at all naturally form a kind. Some inductive inferences are reasonable, and others are not. If there is a paradigm inductive inference at all then this is remarkably different from that for deductive inferences. Perhaps the paradigm inductive inference would be one with a theoretical conclusion that it treats as evidenced by some empirical facts, and yet it explicitly cites in support of the conclusion only those empirical facts. The inference is not an enthymeme for there is no one unique way to fill in the blanks as it were, or thus make explicit all the considerations on the basis of which you can discern, logically, that the empirical facts in question speak to the truth of the conclusion in question. Rather than being an enthymeme, the inference somehow engages the totality of what the inference-maker herself, or herself and her community, receive as knowledge. For this reason, the paradigm inductive inference in fact does not possess the form "premise, premise, …, premise / conclusion". We cross wires with the paradigm for deduction when we suppose that it should.

Surrounding the kind of inductive inference I have just described, there is again a wide penumbra. Some inductive inferences are rather more explicit about the ways that the cited evidence logically speaks to the truth of the conclusion. They may actually make explicit that, say, probability considerations of a certain sort are pertinent to the truth of the conclusion, or issues of explanatory power, or consilience, or parsimony of causal assumptions, or whatever. In that case they may look a bit more like an argument. But unless they are outright demonstrative, and therefore better called "deductive" than "inductive", I don't think that they are best thought of quite as arguments. The extra considerations that they make explicit will however probably guide us to underlying deductive inferences that they synthesise. With appropriate detachment assumptions added in, considerations about probabilities, or explanatory power, or consilience, or parsimony of causal assumptions, or whatever, may be seen to be part of a deductive inference from fallible assumptions to the conclusion of the inductive inference.

The perspective I am defending here is anathema to a lot of philosophers who have attempted to provide a formal theory of inductive inference. Some such attempts achieve some real insights in my view, but I think that these are insights precisely because they can be taken over into the picture that I prefer. The forms of reasoning that have been identified by analytic philosophers who have attempted to describe induction I make out as actually after all those of *deductive* features of the reasoning in question. For example, Bayesians seem to me to draw our attention neither to *the* form of inductive reasoning, nor even to *a* form of it, for inductive reasoning has no form; rather Bayesians explore one form among others for certain deductions to take, that help explicate the reasonableness of an inductive inference that is reasonable. The only difficulty with the

Bayesian analysis is that they don't quite complete the deduction, for they fail to include a detachment assumption, by the inclusion of which the conclusion follows simpliciter, no probabilities in sight.

Imagine some inductive inference and that in light of a Bayesian consideration it is X% probable that its conclusion is true. Then I expect that the following deductive inference partly underlies the reasonableness of the inductive inference in question. From the assumption, obviously fallible but clearly the more reasonable the closer X is to 1, that what is X% probable in the present case also actually happens in the present case to be true, together with the Bayesian consideration, validly infer the conclusion of the inductive inference. In this deduction, our (X%) reasonable but fallible further assumption detaches the conclusion itself of the inductive inference from the Bayesian consideration concerning that conclusion.

The Bayesian consideration itself cannot have been conducted in a theoretical vacuum, so if it is deployed, perhaps tacitly, within the inductive inference, then that inductive inference employs, perhaps tacitly, a host of further background theoretical assumptions that will clearly be fallible. I think that neither the Bayesian consideration on its own, nor the deduction just mentioned that can be formed on its basis, are identical with the inductive inference however. We will surely find under further analysis that there are other reasons besides this one to think that the inductive inference is reasonable.

This is why I contend that, when we analyse (with or without consideration of probabilities) what is reasonable about an inductive inference that is reasonable, we inevitably make it out as synthesising indefinitely many deductive inferences to its conclusion. These deductive inferences all employ background assumptions that are theoretical or otherwise fallible, so they all are fallible, and that is why our analysis confirms that the inductive inference itself was fallible, as will surely be the case.

I say the same about other supposed analyses of inductive reasoning. If the approach in question achieves anything, what it illuminates is best considered in light of a detachment rule. That allows us to see it as part of a background of deductive inference-making. In other words, what one makes out under analysis as reasons why a given inductive inference works are various *deductive* inferences, from admittedly fallible assumptions, to the same conclusion. It is true that, to the extent that such an analysis is at all penetrating, it inevitably elucidates much that in the original inference had been inexplicit and only intuitively gathered. Yet if, as an analysis of the warrant for the given inference, it is any good, it will not be arbitrary in how it does this. Rather it will be sensitive to the epistemic situation of the inference maker, and will address itself only to considerations that were, perhaps merely tacitly, within the inference maker's rational purview. In short, the reasons why a given inductive inference is reasonable, if it is, are bound to be special to the particular subject matter of the inference and to the total epistemic situation of the inference-maker.

When an inference is unreasonable, how this tells against the rationality of the inference-maker depends on whether the inference is deductive or inductive. If the inference is deductive then its unreasonableness tells that the inference-maker failed to be logically consistent in this or that specific respect. That is to say, for faulty deductions, the fault boils down to an inconsistency and only to that extent implicates the broader notion of incoherence. On the other hand, if the inference is inductive, then the problem within the inference-maker is much more general. He failed to get himself together quite generally, that is, to draw into a suitably harmonious relation really very many features all at once of his total epistemic situation. Here, a real problem of incoherence, as opposed to mere inconsistency, is met.

When we analyse the latter kind of failing, we generate lots to say. The more thorough we are in our analysis the more we have to say. When we analyse the former kind of failing, we of course end up with just the one thing to say. The inference-maker went wrong in a quite specific respect.

Similarly if an inference is reasonable, how this tells in favour of the rationality of the inference-maker depends on whether the inference is deductive or inductive. If the inference is deductive then its reasonableness tells that the inference maker achieved a certain specific triumph logically. If the inference is inductive then its reasonableness tells of an embracing kind of harmony or coherence. For example, the question how Wegener induced from empirical evidence that the continents have moved concerns an embracing kind of harmony or coherence. The question how present-day geologists induce from empirical evidence that the continents have moved concerns a harmony or coherence that is still more embracing. To analyse this embracing kind of harmony or coherence sympathetically would begin to bring out various specific logical triumphs with which it is associated, that is, various deductive features that it tacitly has. But while such analysis has much to reveal, it can never reveal all the pertinent, tacit such features. In my §3 I barely scratched the surface of what can be said.

Notably, what analysis reveals as reasons for an induction are arguments for its conclusion that are deductive in form. The various measurement inferences that Wegener brings together to induce that the continents have moved are all deductive inferences. For example, it is by a deduction that Wegener infers from the identity in constitution of two far-flung moraines that they and the continents underneath them once touched. His inference is not of the explanandum (that the moraines are alike) from the explanans (that the continents have moved). Rather it is of the explanans (that the continents have moved) from one of its own explananda (that the moraines are alike). The deduction of a theoretical explanans from one of its own explananda is possible because of help from background theoretical assumptions, such as that no two moraines will be exactly alike in constitution unless they were formed by the same glacier.

On the picture that I defend, inductive inferences are not arguments, and there is no form for an inductive inference to have.□ An analysis of form is certainly called for, but this will concern the underlying deductions.□They are arguments, and they have forms.□

The formal analysis of induction is on my picture misconceived. (But some who have attempted it have incidentally generated some useful insights, insights that I gladly commandeer.)

From all the above we see why there is no single overarching issue about the justifiability of induction. Analysis reveals that an induction synthesises innumerable deductions, all of them fallible, and most or all of them tacit. Thus each induction is justified, if at all, in precisely its own way, which we can illuminate to whatever extent our analytical penetration allows through a discussion that is in the end entirely about deductive inferences. I have called this picture a "disappearance theory" of induction.