INTRODUCTION

We have now seen that any philosophical account of laws needs a good deal in the way of metaphysics to do justice to the concept at all. We have also seen that, as a result, any such account founders on the two fundamental problems of identification and of inference. The extant accounts come to grief additionally in their attempts even to meet the most basic criteria relating to science and explanation. Their promises have all proved empty.

But there are still those traditional arguments, which conclude first that there *must be* laws of nature, and secondly, that we *must believe* that there are such laws. In modern terms, the threats are these: without laws of nature we can make no sense of science and its achievement, nor of rational expectation, and must succumb inevitably to the despair of scepticism.

In this Part I shall answer these epistemological arguments. More constructively, I shall propose a programme for epistemology and for philosophy of science which will allow them to flourish in the absence of laws of nature or belief therein.

6.
Inference to the Best Explanation: Salvation by Laws?

As a man of science you're bound to accept the working hypothesis that explains the facts most plausibly.

The Arch-Vicar of Belial, to Dr Poole, in Aldous Huxley, *Ape and Essence*.[1]

The inference from the phenomena that puzzle us, to their best explanation, appears to have our instinctive assent. We see putative examples of it, in science and philosophy no less than in ordinary life and in literature.

It is exactly this pattern of inference, to the best explanation offered, that philosophers have drawn upon to claim confirmation of laws. They support this appeal in two ways: by pointing to the failures of traditional ideas of induction and by arguing that this inference pattern is the true rock on which epistemology must build. After examining their reasons, I shall argue instead that they would build on shifting sands. As long as the pattern of Inference to the Best Explanation—henceforth, IBE—is left vague, it seems to fit much rational activity. But when we scrutinize its credentials, we find it seriously wanting.[2] (For those more interested in IBE itself than in its connection with laws, sections 2 and 3 may be skipped.)

If both induction and IBE fail as rational basis for opinion and expectation of the future, traditional epistemology is indeed in serious difficulty. But rather than proclaim the death of epistemology, and submit either to an irenic relativism or to sceptical despair, I shall try to show in the next chapter that a new epistemology has been quietly growing within the ruins of the old (as well as show that it issues in a still more drastic critique of IBE).

I. ON THE FAILURES OF INDUCTION[3]

One contention, common to many writers, is that without some such concept as laws of nature, we can make no sense of rational expectation of the future. I have earlier presented this point in what I take to be its primordial form: if anyone says that there is no reason for the observed regularities, then he can have no reason to expect them to continue.

This assertion clearly denies the cogency of induction in a narrow sense (belief based on straight extrapolation from the data) while it holds out the hope of induction in a very broad sense (rationally formed expectation of the future). Let us here use the term 'induction' everywhere in its narrow sense: the procedure whose independent rationality friends of laws tend to deny. I may as well add at once that I agree with them on the critical point. My discussion will aim to underline their legitimate objections, but to show simultaneously how their critique is misdirected, and where it rests on dubitable premisses of their own.

Here is the ideal of induction: of a rule of calculation, that extrapolates from particular data to general (or at least ampliative) conclusions. Parts of the ideal are (*a*) that it is a *rule*, (*b*) that it is *rationally compelling*, and (*c*) that it is *objective* in the sense of being independent of the historical or psychological context in which the data appear, and finally, (*d*) that it is *ampliative*. If this ideal is correct, then support of general conclusions by the data is able to guide our opinion, without recourse to anything outside the data—such as laws, necessities, universals, or what have you.

Critique of this ideal is made no easier by the fact that this rule of induction does not exist. The rule was indeed baptized—presumably after conception, but before it was ever born. Sketches of rules of this sort have been presented, with a good deal of hand-waving, but none has ever been seriously advocated for long.[4] Every generation of philosophers, beginning with Aristotle, has seen that mere numerical extrapolation in any specific form, cannot be the rule described in our ideal. Criticisms brought forward in this century however, exhibit difficulties to plague every possible realization of the ideal. If the reader is already convinced of the inadequacy of induction in the narrow sense, there is no need to read the rest of this section.

What is extrapolation? Example of the alien die

Consider a die, which is to be tossed ten times, and the hypothesis that all ten tosses will come up *ace*. Let the evidence so far be that it has come up *ace* for the first seven tosses. Now, how could the rule of induction relate these data to the hypothesis? Should it tell us that all ten will be like the first seven? Rules of extrapolation can't be expected to do well if some relevant evidence is left unstated, and we do have other information about human dice. So suppose we found this die on an alien planet, and 'ace' is the name we give to one of the six sides of this geometric cube of unknown composition.

Of course, reader, you are still unwilling to suggest that the rule should tell us to infer that all tosses (or equivalently, the last three) will come up *ace*. After all, in this situation, *you* would not infer that. But perhaps the rule should be sophisticated beyond anything Bacon and Newton, the great advocates of this ideal, could imagine. Let it tell us the probability of the hypothesis, bestowed on it by the data. Then it could be asserted that a rational person must follow the rule of induction, in the sense that it provides him with the probability that the hypothesis is true, given the evidence.

This suggestion marks quite a shift, because it takes us from induction as extrapolation from mere numbers to something much more general. But every discussion of induction is forced to this. Suppose, for example, that instead we try to maintain the rule in as simple a form as possible, with as one corollary: if all instances have been favourable, and you have no other evidence, then believe that the next instance will be favourable as well. You will immediately insist, surely, that 'believe' must be qualified here, if you are ever to follow it. Believe with what confidence? Believe to what degree? Are a hundred instances not better evidence than ten, even if all have been favourable? Any such reaction replaces the simple rule with a more sophisticated one, of the order of probability assignments. Of course we should hasten to accept all worthy suggestions for improvement of the rule, rather than insist on beating a dead horse.

In both forms, the same problem about induction appears very clearly. It is that, being a *rule* it must have certain structural features—and as a result, its extrapolation from any data will be heavily influenced by what it does with small increases in data. But how it does that, must be either uninformative or arbitrary. That is the dilemma this ideal always foundered on.

To illustrate this, begin with the naïve 'straight rule': believe that the ratio of A to B overall equals the ratio observed so far. That tells us to believe that the sun will always rise. Unfortunately, it also tells us to expect all *aces* as soon as we have seen a single toss of the alien die, if it came up *ace*. We can't very well suggest that another ratio is any more plausible. So we must fiddle with the confidence: always believe what this rule tells you, but believe it weakly after one toss and strongly after many tosses.[5] Now what is needed is an exact prescription of what beliefs, degrees of belief, or confidence I should have at the outset; and an auxiliary rule about how this should change with the outcome of each new toss. The former could be perfect neutrality of opinion. (I do not assume it must be a precise subjective probability or anything like it.) But then the auxiliary rule must still say exactly how much non-neutrality there should be after one toss; and indeed, after $n + 1$ tosses (as a function of n, the previous outcomes, and the new outcome). Now you can look at the numbers and ratios as much as you like, but they will give you no clue at all to this auxiliary rule for massaging your confidence in the observed ratio. Myriads, continua, of such functions exist, and however little they diverge in the small, they lead to widely different consequences down the line. You can try to remain a little neutral among them: the more you do so, the *less arbitrary* will your rule of induction be, but also the *less informative*. Now, next problem: try and formulate a measure of balance between arbitrariness and informativeness. There you will again find a continuum of functions to choose from, and you will again confront the dilemma presented by the spectrum from capricious choice to trivializing neutrality.[6]

There are many other problems with the ideal of induction, even if this is (as I think) its fundamental flaw.[7] As to the empiricists who followed this banner *sans* device, their hope was placed in an empty promise. But does there indeed lie a better hope in the mobilization of laws to found rational expectation, as Dretske, Armstrong, and Tooley contend?

2. DRETSKE ON THE REMARKABLE CONFIRMATION OF LAWS

The scheme we inspected above might be called simple or bare induction; many have been the proposals to replace it by more complex schemes. One such is Dretske's proposal that we take note of the (supposed) remarkable tendency of laws to become well confirmed on the basis of very little evidence. He uses exactly the reflections and types of examples exhibited above, to sketch a rival picture of how we can rationally go beyond the evidence.

It appears, at first sight, that laws are beyond our epistemic grasp. The sun has risen every day; this appears to confirm the universal generalization that it always has and always will. But if we speculated that this was a matter of law, we would be asserting more: something beyond and in addition to the universal statement. The conforming instances support the generalization, but surely they do not support anything beyond that? Dretske tells us that this apparent problem is a pseudo-problem, resting on a mistaken empiricist epistemology. In fact, he claims, it is quite the other way round. This sort of evidence, of positive instances, does not at all support the universal statement, if taken in isolation. But it does support the hypothesis that it is a law that the phenomenon always occurs in that same way.

Dretske calls this conclusion 'mildly paradoxical' ('Laws of Nature', p. 267), but it seems more than that. Surely if it is a law that A then it is also true that A; hence I can become no less confident that it is so than that it is a law. The air of paradox is perhaps removed, if we take Dretske to be attacking the conception of evidential support that was implicit in the proffered argument. This conception appears to be the old ideal of purely numerical induction.

To tackle also more sophisticated epistemic schemes, Dretske makes the preliminary point that raising our probabilities may not amount to real confirmation.[8] Let us use the alien die, introduced in the preceding section, to illustrate his point. Suppose I begin with the initial assumption that the die is fair. Then my initial probability that it will come up *ace* all ten tosses, is very low—$(\frac{1}{6})^{10}$ which equals about six in a thousand million. After I have seen seven *aces* come up in in a row, while maintaining this assumption, the probability that the last three will come up, is still the same as it was: $(\frac{1}{6})^3$. But this is now also the new probability of the proposition that all ten come up *ace*. The probability of that proposition has therefore become 6^7 times—about 300 000 times—higher than it was. Our probability for the universal statement has increased dramatically—but we are in no better position to predict what comes next! Using a similar example about coin tossing, Dretske writes:

> But this, of course, isn't confirmation. Confirmation is not simply raising the probability that a hypothesis is true, it is raising the probability that the unexamined cases resemble (in the relevant respects) the examined cases. It is *this* probability that must be raised if genuine confirmation is to occur (and if a confirmed hypothesis is to be useful in *prediction*), and it is precisely this probability that is left unaffected by the instantial 'evidence' in the above examples. ('Laws of Nature', p. 258)

> The only way we can get a purchase on the unexamined cases is to introduce a hypothesis which, while *explaining* the data we already have, *implies* something about the data we do not have. (ibid. 259)

Let us criticize the example and its discussion immediately.

The moral is not correctly drawn, because it is only on the supposition of one explanatory hypothesis (e.g. fairness, or any other sort of *bias*) that the data can raise the probability of the universal statement without raising that of the remaining instances. If instead I profess some measure of ignorance about the bias of the die, then that ignorance becomes modified by the initial data, and my opinion about the unexamined cases changes right along with it. For example, if I had thought that the die was *either* fair *or* perfectly biased in favour of *ace*, then after seven *aces* I would have favoured the latter hypothesis considerably! I would accordingly think it more likely then than I did before that the last three would be *aces* too. So Dretske has generalized upon a special case.

The second point that had better be noticed is that these effects would appear in the same way if our background beliefs had nothing lawlike about them, as long as they relate to the instances in the same way. For consider another example. I am told that the ten coins I am about to be shown came either from Peter's pocket or from Paul's; that Peter's contained ten dimes and fifty nickels, while Paul's contained sixty dimes. The first seven to be put before me are dimes. Obviously, on the supposition that they all come from Peter's pocket, the hypothesis that all ten will be dimes has increased to the constant probability that the last three will be, namely $(\frac{1}{6})^3$. Without this supposition, but with the background belief that they are equally likely to come from Paul's as from Peter's, the probability of the last three being dimes has also

increased, however. This example parallels the previous one perfectly, although here no laws of nature are involved at all. The non-lawlike hypothesis about pockets has the same effect. None of this has anything to do with the explanatory power that may or may not reside in the lawlikeness of background opinion, but only with what that opinion is (as expressed in terms of my personal probability).

Of course, someone else might be happy to say that the hypothesis, that all the coins come from Peter's pocket, explained why the first seven were dimes. But Dretske's contention appears to be that to be explanatory, the hypothesis has to be about something special, like laws or similar unordinary facts.

Dretske's alternative proposal

Suppose, however, that we do agree to this idea about what explanation requires. What rival to primitive numerical induction does Dretske want to propose? He proposes that we follow a rule of IBE.

If the first seven tosses yield *ace* this is best explained by the die's having a perfect bias in favour of *ace* (let us say). Should we now at once accept this hypothesis? But it was the best explanation already when we had just seen the first toss yield *ace* and for the same reason. Should I therefore have accepted the hypothesis of extreme bias already after one *ace*? Obviously not; so I am construing the proposal too naïvely. It cannot be that I'm simply to infer the truth of the best explanation. Rather, the all-or-nothing model of jumping from agnosticism to full belief must again be modified to accommodate, and trade in, degrees of belief or confidence. That conclusion is also evident in Dretske's discussion, though he talks of confidence rather than of probability:

> laws are the *sort* of thing that can become well established prior to an exhaustive enumeration of the instances to which they apply. This, of course, is what gives laws their predictive utility. Our confidence in them increases at a much more rapid rate than does the ratio of favourable examined cases to total number of cases. (ibid. 256)

This is not as easy to construe as it looks! Let us be careful, and see what meanings this passage can and cannot bear.

First, it cannot mean that our confidence in a proposition increases rapidly with accumulating evidence, if we know (or believe) that it is a law. For then our confidence is already at a maximum.

Second, if it really were a law that all A are B, would that by itself make our confidence in it increase especially fast, regardless of what we believe at the outset? Surely not—the law of gravity may make things fall, but can't make people believe that things will fall.

Third, might Dretske mean that, in response to the same evidence, my confidence in the proposition *It is a law that all A are B* increases more rapidly than my confidence in *All A are B*? That could be. But since the former is supposed to entail the latter, it will catch up, and from there on must inevitably drive the latter ahead. For example, it seems likely to be a law that all radium decays, I must then regard it at least as likely that all radium does decay.

Besides these three possible meanings which Dretske cannot intend, how else could the assertions be construed? Perhaps the fault lies with our lack of imagination. Possibly Dretske is pointing to a rule, as yet unknown, which will make or revise our probabilities so as to give *bonus* marks to the hypotheses that explain observed phenomena. Then, if hypotheses to the effect that there exists a law are especially explanatory, they may get an especially high *bonus*. This suggestion cannot be dismissed, nor discussed probatively at the merely qualitative level. I shall discuss it in the next chapter, and argue that there cannot exist any such probabilistic rule of Inference to the Best Explanation, on pain of incoherence. We will also find there that other, more precise construals of Dretske's dictum leave it equally false.

3. ARMSTRONG'S JUSTIFICATION OF INDUCTION

At the beginning of his book, Armstrong wrote 'There is one truly eccentric view, brought to my attention by Peter Forrest. . . . This is the view that, although there are regularities in the world, there are no laws of nature.' Armstrong's response follows at once: 'Such a view, however, will have to face the question what good reason we can have to think that the world is regular. It will have to face The Problem of Induction. It will be argued . . . that [no such view] can escape inductive scepticism' (*What is a Law of Nature?*, p. 5). He makes good his promise in a later chapter, by a real *tour de force*. For there Armstrong purports not only to prove that belief in laws is needed, to avoid the sceptic's slough of despond, but to present us with a justification of induction—by an argument which does not depend on what the rules of induction are! I shall now analyse this carefully to show what he assumes along the way.[9] What this analysis will show, among other things, is that Armstrong's own argument *relies* on a previously assumed rule of Inference to the Best Explanation, and advances no independent support for it.

Armstrong begins with the explicit premiss (call it P_1) that 'ordinary inductive inference, ordinary inference from the observed to the unobserved, is . . . a rational form of inference'. On questions about what that form is, what rules may be being followed, he confesses himself largely agnostic. He defends the premiss along the lines of Moore, common sense against scepticism, saying that this premiss is part of the bedrock of our beliefs, indeed, that it 'has claims to be our most basic belief of all' (p. 54).

This premiss (P_1) is theoretically loaded despite the accompanying agnosticism on questions of form. It is undoubtedly true that we have expectations about the future, and opinions about the unobserved. It does not follow that we are engaged in ampliation—let alone some sort of ampliative *inference*, i.e. ampliation in accordance with rules. Perhaps we amend our opinions (*a*) by purely logical adjustment to the deliverances of new experience ('conditionalization', for example) and/or (*b*), some unpredictable free enterprise in the formation of new opinions, within certain limits required by rationality. The distance Armstrong slides here stretches from the Moorean common sense that we form rational expectations, to the philosophical modelling of this activity as *inference*.

Besides the explicit premiss, therefore, we have found as further premiss the statement that we believe the initial premiss, and that we either know or rationally believe it to be true. This extra premiss (call it P_2) is needed to understand the subsequent argument. Yet it is seriously questionable.

Suppose now that we do engage in some form of ampliative inference, which we believe to be rational. At this point in the argument, we need not yet know what that form is (one form that fits all ampliative inference is 'P; therefore Q', but the 'ordinary inductive inference' presumably includes much less than everything fitting that form)—so let us call this form F.

Here follows the *first sub-argument*. Its conclusion is that it is a necessary truth that induction (i.e. in our present terms, ampliative inference of form F) is rational. This argument is based on P_2 rather than on P_1 and is an interesting variant on Peirce's argument for the reality of laws. P_2 says that we know or rationally believe induction to be rational. But that implies (via a premiss which I shall not number) that our belief that induction is rational must have a justification. That justification cannot be by induction or it would be circular. Nor can it be by deduction, since the relevant statement (i.e. P_1) is not a logical truth and any premiss from which P_1 could be deduced would face the same question as we have for P_1, thus leading to a regress. The only possible justification is therefore a claim to knowledge or rational belief not based on any sort of demonstration. That is a tenable claim only for a statement claimed to be known a priori. But (again via a premiss I shan't number) only necessary truths can be known a priori. Therefore P_1 is a necessary truth.

I am not sure that to be rational a belief must have a justification reaching back all the way to a priori truths. But if we allow, say, a priori truths plus the evidence 'of my own eyes', as basis for justification, the case for P_1 will not be significantly different. So Armstrong has 'established' that if induction is known to be rational, this must be a case of a priori knowledge. And if only necessary truths can be known or rationally believed without the sort of justification that P_1 is denied, then P_2 implies that P_1 is a necessary truth.

We come now to the *second sub-argument*. The conclusion we have reached is this: it is a necessary truth that ampliative induction of form F is rational. This fact (call it P_3) Armstrong insists, must be given an explanation. What F is will now finally make a difference. He proceeds as follows: he makes a proposal for what F is, demonstrates how P_3 can be true on the basis of this proposal, and then notes that the demonstration would fail if laws said nothing more than mere statements of regularity.

The proposal is that ampliative inference of form F is inference from the evidence, to laws that explain the evidence (call this P_4). He adds that this procedure is an instance of Inference to the Best Explanation, that this sort of inference (IBE) is rational (P_5), and

that it is analytic (true by virtue of the meanings of the words) that IBE is rational (P_6). If we add that analytic statements are necessary truths, the explanation of (P_3) is complete. The footnote to be added about any 'regularity' view of laws is that it would make (P_4) false, because according to Armstrong regularities, unlike laws, do not *explain* the evidence which they fit or entail.

Note well that (P_4), (P_5), and (P_6) are not premisses of the overall argument. They are premisses of a sub-argument, whose correctness—once noted—is all that is asserted. Because it is correct, it gives us an explanation of (P_3). Since no other explanation of (P_3) is available, it is supposedly rational to believe the explanation offered, and hence the premisses on which it rests. *Here*, in this ultimate stage of the argument, we see a step made by means of inference to what explains. So IBE does function as premiss of the overall argument as well. What independent support could any of these premisses receive?

The defence of (P_4) must be that laws explain the phenomena which they fit or entail, and that laws provide the best among the explanations that can be given for such phenomena. This second part of the defence comes in a very cavalier little paragraph:

> It could be still wondered whether an appeal to laws is really the *best* explanation of [the phenomena]. To that we can reply with a challenge 'Produce a better, or equally good, explanation'. Perhaps the challenge can be met. We simply wait and see. (p. 59)

Would it be enough, to meet this challenge, to present some cases where the best available explanation of some phenomena does not consist in deriving them from laws? If so, there is enough literature for Armstrong to confront now; he need not wait.

The defence of the next premiss, (P_5)—namely, that IBE is rational—consists simply in the last premiss, (P_6): it is analytic, due to the meaning of the word 'rational', that IBE is rational. This conviction about IBE appears not only here, but throughout the book, and not surprisingly: IBE is the engine that drives Armstrong's metaphysical enterprise. It provides his view of science (p. 6: 'We may make an "inference to the best explanation from the predictive success of contemporary scientific theory to the conclusion that such theory mirrors at least some of the laws of nature . . ."'). He also regards IBE as being first of all a form of inference to be found pervasively in science and in ordinary life (p. 98: 'But I take it that inference to a good, with luck the best, explanation has force *even* in the sphere of metaphysical analysis'—(my italics)). To support (P_6) he does not see the need for more than rhetorical questions: 'If making such an inference is not rational, what is?' (p. 53); 'To infer to the best explanation is part of what it is to be rational. If that is not rational, what is?' (p. 59).

In sum, therefore, Armstrong has reached powerful conclusions on the basis of an assumption which is supported solely by a challenge to those who would doubt it.

In the next two sections I hope to meet Armstrong's challenge. I shall argue that inference to the best explanation cannot be a recipe for rational change of opinion. And then I shall try to answer the question, 'If that is not rational, what is?'

4. WHY I DO NOT BELIEVE IN INFERENCE TO THE BEST EXPLANATION[10]

There are many charges to be laid against the epistemological scheme of Inference to the Best Explanation. One is that it pretends to be something other than it is. Another is that it is supported by bad arguments. A third is that it conflicts with other forms of change of opinion, that we accept as rational.

Still, the verdict I shall urge is a gentle one. Someone who comes to hold a belief because he found it explanatory, is not *thereby* irrational.[11] He becomes irrational, however, if he adopts it as a rule to do so, and even more if he regards us as rationally compelled by it. The argument for this conclusion will be begun here and concluded in the next chapter.

What IBE really is

Inference to the Best Explanation is not what it pretends to be, if it pretends to fulfil the ideal of induction. As such its purport is to be a rule to form warranted new beliefs on the basis of the evidence, the evidence alone, in a purely objective manner. It purports to do this on the basis of an evaluation of hypotheses with respect to how well they explain the evidence, where explanation again is an objective relation between hypothesis and evidence alone.

It cannot be *that* for it is a rule that only selects the best among the historically given hypotheses. We can watch no contest of the theories we have so painfully struggled to formulate, with those no one has proposed. So our selection may well be the best of a bad lot. To believe is *at least* to consider more likely to be true, than not. So to believe the best explanation requires more than an evaluation of the given hypothesis. It requires a step beyond the comparative judgment that this hypothesis is better than its actual rivals. While the comparative judgment is indeed a 'weighing (in the light of) the evidence', the extra step—let us call it the ampliative step—is not. For me to take it that the best of set X will be more likely to be true than not, requires a prior belief that the truth is already more likely to be found in X, than not.

There are three possible reactions to this, each of which argues that IBE must be allowed *nevertheless* to play the role of leading to a new belief extrapolated from one's evidence. Clearly any such reaction must focus on the ampliative step, because the above objection is independent of the method of evaluation (of explanatoriness) that is used.

Reaction 1: Privilege

The first consists in a claim of privilege for our genius. Its idea is to glory in the belief that we are by nature predisposed to hit on the right range of hypotheses.[12]

We recognize here the medieval metaphysical principle of *adequatio mentis a rei*. Contemporary readers will not be happy to accept it as such, I think, and would hope for a justification. Such a justification could take two forms, allied respectively with naturalism and rationalism in epistemology.

The naturalistic response bases the conclusion on the fact of our adaptation to nature, our evolutionary success which must be due to a certain fitness. But in this particular case, the conclusion will not follow without a hypothesis of pre-adaptation, contrary to what is allowed by Darwinism.[13] The jungle red in tooth and claw does not select for internal virtues—not even ones that could increase the chance of adaptation or even survival beyond the short run. Our new theories cannot be more likely to be true, merely given that we were the ones to think of them and we have characteristics selected for in the past, because the success at issue is success in the future. The moths in industrial England became dark, not because they began to have more dark offspring but because the light ones were more vulnerable.

The rationalist response must be patterned after Descartes's argument for the correspondence of ideas to reality. Alvin Plantinga has suggested such a reason for privilege: given other beliefs about God, such as that we are made in his image, it is only reasonable to believe that we are specially adapted to hit on the truth when we come up with our (admittedly limited) guesses at explanation. Plantinga applies this even to belief in propositions and other abstract entities. But it takes more than a generally agreed concept of God to get this far. For even if he created us naturally able to perceive the truth about what is important for us in his eyes (perhaps to discern love from lust, or charity from hypocrisy, in ourselves), this may not extend to speculations about demons, quarks, or universals. *Privilege* is consistent, but seems incapable of either naturalistic or rationalist support.

Although it does not count as an argument, I should also point out that Privilege is entirely at odds with Empiricism. By this I mean the position that experience is our one and only source of information. Clearly this leaves open a great deal—experience may be very rich in its possible varieties; on the other hand, the information it brings us may not come as if in the voice of an angel, but in dubious and defeasible form. However that may be, the position sets one clear limit: if we do have innate or instinctive or inborn expectations, we'll be just lucky if they lead us aright, and not like lemmings into the sea. However basic or natural our inclination may be toward, for example, more satisfying explanations, that inclination itself cannot be relevant information about their content's truth.

Reaction 2: Force majeure

The *second* reaction pleads *force majeure*: it is to try and provide arguments to the effect that we *must* choose among the historically given significant hypotheses. To guide this choice is the task of any rule of right reason. In other words, it is not because we have special beliefs (such as that it will be a good thing to choose from a certain batch of hypotheses), but because we must choose from that batch, that we make the choice.

The *force majeure* reaction is, I think, doomed to fail. Circumstances may force us to act on the best alternative open to us.

They cannot force us to believe that it is, *ipso facto*, a good alternative.

Perhaps it will be objected that the action reveals the belief, because the two are logically connected. And in a general way that is certainly true. But it is exactly in situations of forced choice that action reveals very little about belief. Think of William James's example of a walker in the mountains who has the choice of jumping over a crevasse, or remaining for the night. Suppose that both a fall and exposure mean almost certain death. The prize is equal too: life itself. She jumps. Does this reveal that it was very likely to her that she would get across? Not at all, for even a very low chance would be reasonable to take in this case. Or if the Princess must open one of N doors, the Tigers skulking behind all but one, we can only conclude about the door she opened, that it seemed to her no less probable than $1/N$ that it would lead to freedom. And if there might be a Tiger behind each, and she is forced to choose nevertheless, we cannot even conclude that.

In the case of science we certainly observe theory choice. But just what belief is revealed there? Let us look carefully at the practice, and see what belief it entails, if any.

Scientists designing a research programme, bet their career and life's satisfactions on certain theoretical directions and experimental innovations. Here they are forced to choose between historically given theoretical bases. They are forced by their own decision to be scientists, to opt for the best *available* theory, by their own light. What beliefs are involved in this, can be gauged to some extent from their goals and values. Does this scientist feel that his life will have been wasted if he has spent it working on a false theory? Then he must feel that Descartes's and Newton's lives were wasted. Or does he feel that his life will have been worth while if he has contributed to progress of science, even if the contribution consisted in showing the limits and inadequacies of the theories he began with and the discovery of some new phenomenon that every future science must save? In the latter case his choice between theories, as basis for research, does not reveal any tendency to belief in their truth.

Reaction 3: Retrenchment

The *third* reaction is to *retrench*: 'Inference to the Best Explanation' was a misnomer, and the rule properly understood leads to a revision of judgement much more modest than inference to the truth of the favoured hypothesis. The charge should be that I have construed the rule of inference to the best explanation too naïvely. Despite its name, it is not the rule to infer the truth of the best available explanation. That is only a code for the real rule, which is to allocate our personal probabilities with due respect to explanation. Explanatory power is a mark of truth, not infallible, but a characteristic symptom.

This retrenchment can take two forms. The *first form* is that the special features which make for explanation among empirically unrefuted theories, make them (more) likely to be true.[14] The *second form* is that the notion of rationality itself requires these features to function as relevant factors in the rules for rational response to the evidence. I will take up both forms—the first in the remainder of this section, and the second in the next chapter. Let us note beforehand that the first must lean on *intrinsic explanatoriness*, which can be discerned prior to empirical observations, and the second specifically on *explanatory success* after the observational results come in. What the criteria are for either, we shall leave up to the retrencher.

Retrenchment, form 1

Is the best explanation we have, likely to be true? Here is my argument to the contrary.

I believe, and so do you, that there are many theories, perhaps never yet formulated but in accordance with all evidence so far, which explain at least as well as the best we have now. Since these theories can disagree in so many ways about statements that go beyond our evidence to date, it is clear that most of them by far must be false. I know nothing about our best explanation, relevant to its truth-value, except that it belongs to this class. So I must treat it as a random member of this class, most of which is false. Hence it must seem very improbable to me that it is true.

You may challenge this in two ways. You may say that we do have further knowledge of our own best explanation, relevant to its truth-value, beyond how well it explains. I'm afraid that this will bring you back to the reaction of Privilege, to glory in the assumption of our natural or historical superiority. Or you may say that I have construed the reference class too broadly. That is fair. The class of rivals to be considered should be on a par with

our own, in ways that could affect proportions of truth. So we cannot include in it two theories, each having the same disagreement with ours at point X, but then disagreeing with each other at point Y, on which ours has nothing to say. So for each statement ours makes, beyond the evidence we have already, we can include only one theory disagreeing with that statement, for every way to disagree with it. But of course, there is only one true way to agree or disagree at this point. So the conclusion, that most of this class is false, still stands.

David Armstrong, replying to a version of this argument, writes 'I take it that van Fraassen is having a bit of fun here.'[15] Yes, I had better own up immediately: I think I know what is wrong with the above argument. But my diagnosis is part and parcel of an approach to epistemology (to be explained in the next chapter), in which rules like IBE have no status (*qua* rules) at all. As a critique of IBE, on its own ground, the above argument stands.

One suspicious feature of my above argument is that it needed no premisss about what features exactly do make a hypothesis a good explanation. Let us consider a contrary argument offered by J. J. C. Smart, which does focus on one such feature, simplicity. I think that both argument and counterargument will be rather typical of how any such debate could go (for any choice of explanatory feature). Smart begins as follows:

> My argument depends on giving a non-negligible a priori probability to the proposition that the universe is simple. . . .
>
> Let p = the observational facts are as if there are electrons, etc.
> q = the universe is simple
> r = there really are electrons, etc.
>
> We can agree . . . that $P(p) > P(pr)$, as of course we have to! But I want to say $P(pr/q) > P(p\bar{r}/q)$.[16]

That is, the probability *on the supposition that the universe is simple*, is greater for our best explanation which entails that the phenomena will continue as before, than it is for a denial of that explanation which agrees on those phenomena.

The reason for this judgement, on Smart's part, must hinge on a connection between explanation and simplicity. So it does. It is exactly because the explaining hypothesis is simple (as it must be, to qualify as explanation) that the supposed simplicity of the universe makes the hypothesis more probable than its denial (under the further supposition that the explanation is right about all the phenomena, even those to come). And it seems at first sight plausible to say that a supposed simplicity in nature, will make simple theories more likely to be true.

But this plausibility derives from an equivocation. If the simplicity of the universe can be made into a concrete notion by specifying objective structural features that make for simplicity, then I can see how Smart may have arrived at the opinion that the universe is (probably) simple. For there can be evidence for any objective structural feature. But if the universe's simplicity means the relational property, that it lends itself to manageable description by us (given our limitations and capacities) I cannot see that. The successes we've had are all successes among the descriptions we could give of nearby parts of the universe, and of the sort which our descriptive abilities allow. Suppose it is true that the frog can distinguish only the grossest differences between objects at rest, but can notice even small moving objects. Then his success in catching insects flying by is no index of how many potential prey and potential enemies sit there quietly watching him.

Could simplicity in the first sense, which we might have reason to surmise, make more probable the simpler-for-us among the accounts we can give? That depends in part on how much simplicity of theories has to do with simplicity of the world as described by those theories. But suppose there is an intimate connection—unlike, say, in literature where the simplicity and economy of form in poetry does not limit it to simplicity of subject, in comparison with prose. Then still the allocation of probabilities is effectively prevented by a very modest consideration. Simplicity is global. A part of a structure, which is very simple overall, may be exceedingly complex considered in isolation. Here is a simply described set: that of all descendants of Geoffrey of Monmouth, alive today. Now, try to describe it purely in terms of features recognizable today (by geographical location, blood type, what have you) without reference to the past. Considered as part of a historical structure, the set is easily delineated; viewed short-sightedly in the twentieth century it is not. Similarly for the simplicity of the universe as a whole, and those aspects and parts of it on which our sciences focus. The simplicity a situation has in virtue of being part of a simple universe, does not make more likely any simple putative description of it by itself alone.

Retrenchment, form 2

I already raised the possibility of a sophisticated, probabilistic version of the rule briefly above. Combining the ideas of personal probability and living by rules, the new rule of IBE would be a recipe for adjusting our personal probabilities while respecting the *explanatory* (as well as predictive) success of hypotheses. This will be investigated in the next chapter, after we have taken a much closer look at the representation of opinion.

5. WHAT GOOD IS THERE IN THIS RULE?

There must clearly be a solid basis for the appeal and renown that IBE has enjoyed over the years. It is important to ferret this out. But as with any subject, we must carefully separate the inflated claims of philosophical exponents from the grains of common sense which gave those claims their initial appeal. Eventually we must also show that the common-sense part is equally respected in our own account.

If I already believe that the truth about something is likely to be found among certain alternatives, and if I want to choose one of them, then I shall certainly choose the one I consider the best. *That* is a core of common sense which no one will deny.

But how far is it from this common sense to IBE, conceived as cornerstone of epistemology? This rule cannot supply the initial context of belief or opinion within which alone it can become applicable. Therefore it cannot be what 'grounds' rational opinion.

That is only the first point. Next we can see that even if the rule is applicable, we might very well not wish to apply it. Suppose it seems likely to me that one of the first six horses will win the race, and that of these, horse No. 1 is the best. It does not follow at all that I shall then wish to predict that horse No. 1 will win, for this might mean no higher probability than $\frac{1}{5}$ for its winning. Similarly, if I turn away then, and just at the end of the race a great cheer goes up, the best explanation for me of this cheer will be that horse No. 1 has just won. And still I shall be no readier to say that this is what has just happened. If a *force majeure* makes me predict, then I shall indeed say 'No. 1'—but this will certainly not reveal then that I believe it.

A rule which we would often, when it is applicable, prefer not to apply, is not a rule we are following! Common sense will often prevent us from applying IBE; and it could not very well do that if it were the epistemic categorical imperative. Even more important, perhaps, is this: even if the rule is applicable, and we make our choice in accordance with it, we may not be following the rule. This sounds paradoxical, but think: must a choice among hypotheses, even if unconstrained, necessarily be a choice to believe?

In general, the common-sense choice will be in the context of an opinion to the effect that this batch of hypotheses have a balance of certain virtues, and are well fitted to serving our present aims. Likelihood of truth will presumably be among those virtues, or among those aims, but it need not be alone. Informativeness with respect to topics of interest may be another. *When we choose the best, therefore, the choice must be interpreted in terms of the basis for choice.* If likelihood of truth is not the sole basis, then the choice—choice to accept—must not be equated with choice to believe.

How little comfort common sense gives to philosophical fancy! What has happened to extrapolation of general truth from evidence alone? No more than the metaphysician do I think that common sense brings its own clarity with it. So I propose next to discuss the whole subject again in a higher—and more constructive—key. Specifically, I will consider versions of IBE that make full use of the conceptual resources of probability. And yet, its fortunes will not improve thereby.

7.
Towards a New Epistemology

> . . . it is by no means clear that students of the sciences . . . would have any methodology left if abduction is abandoned. If the fact that a theory provides the best available explanation for some important phenomenon is not a justification for believing that the theory is at least approximately true, then it is hard to see how intellectual inquiry could proceed.
>
> Richard Boyd[1]

So far I have only offered a critique of traditional epistemology with its ampliative rules of induction and inference to the best explanation (abduction). But this mainstream does not constitute the only tradition in epistemology. The seventeenth century gave us besides Descartes and Newton also Blaise Pascal, and from his less systematic writings there sprung a stream that in the succeeding three centuries has become a powerful river: the underground epistemology of *probabilism.*

After introducing its basic ideas, I shall show that it leads us into a much more radical and far-reaching critique than we have seen so far. The rule of IBE, and indeed the whole species of ampliative rules, is incoherent. I shall deliver this critique at the outset through a proxy, a foil, one particular sort of probabilist: the orthodox Bayesian. The rigours of *his* views are however also considerably more than his arguments can demonstrate. I will go on then to propose an epistemology that is certainly still probabilist, but offers a reconciliation with traditional epistemology. It does give room to practices of ampliation beyond the evidence. (In the next chapter, we shall see that it also allows us finally to illuminate how objective chance should guide our personal expectation—remember the fundamental question, concerning chance, which previous chapters had to leave unanswered?—and how it may enter the opinion that guides rational decisions.) To end I shall argue that, despite the ominous warnings of the past, our new epistemology is driven into neither sceptical despair, nor feeble relativism, nor metaphysical realism.