The Problem of Induction

JAMES CARGILE

No one doubts that philosophers have discussed at length 'the problem of induction', but it would also be generally recognized that there would be disagreement as to precisely what that problem is. Rather than tackle the formulation problem, I will borrow from a popular text:

> Our existence as well as science itself is based on the principle of induction that tells us to reason from past frequencies to future likelihoods, from the limited known of the past and present to the unknown of the past, present, and future ... But though inductive probability is psychologically inescapable, we have trouble providing a rational justification for it.[1]

We might say, then, that there is such a practice as induction, and a problem associated with it is that of justifying engaging in it. We engage in reasoning from things we know about the past and present to conclusions about the past, present and future. We can't resist doing this but we have trouble finding a rational justification for doing so. This problem suggests a generalization. We engage in reasoning, reaching new conclusions. It would be hard to resist engaging in this practice. How do we provide a rational justification for it?

It seems appropriate to respond to such a question with the observation that some of our reasoning is not justifiable, being badly done. It should be beyond dispute that an attempt to justify all reasoning generically would be hopeless. Less indisputable, though also true, is that the justification of conclusions which are in fact justified is not generic either, and depends on the particular case. Against this, it may be said that there has to be something common to all justified conclusions which are justified, which is what would be sufficient to cite in response to the challenge to justify that conclusion. The inclination to such a view as this may seem to be one which twentieth century philosophy has thoroughly dispelled. But this is not as clear as it might seem, especially when we go back to the subclass of conclusions supposedly marked off by calling them inductive. Some think that these are a special class of

[1] *The Theory of Knowledge, Classic and Contemporary Readings*, Louis P. Pojman (ed.) (Wadsworth, 1993), 430.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

justified conclusions. Others think that they are a special class of conclusions such that those of them which are justified require a generic kind of justification. But one of the most interesting features of this topic is that even philosophers who reject the idea of a generic justification for induction end up offering what, if correct, would have to be just that. This is a phenomenon worth study, which might help explain why induction is such a perennial source of philosophical puzzlement.

Hume asked 'What is the foundation of all conclusions from experience?'[2] His answer was that such conclusions have no foundation in reasoning. 'All inferences from experience ... are effects of custom, not of reasoning.' (p. 57) This might be merely a certain use of 'reasoning' which sounds peculiar to those of us who use the term differently. His saying that 'causes and effects are discoverable, not by reason but by experience' could be taken as a more or less 'reasonable' pronouncement of empiricism. More reasonable: what will happen cannot be deduced from premises solely about what has happened with logical validity. This claim, that there is 'no necessary connection' between what has happened and what will, needs explaining, but it has a basis in truths about necessity.

It is even true that Hume had some reason to think this true observation about necessity was denied by historical rationalists. For they are associated with the idea that laws of nature must be necessary truths. Here there is a danger of misunderstanding. Aristotle means by a necessary truth, a proposition that is true at all times. Even this saying is ambiguous. All true propositions are true at all times, in the proper understanding of these matters. But ignoring that for now, a proper natural law, in the Aristotelian scheme, would be that all A's *always* turn out to be B's, not just up to or during a certain time period. Being timeless in this way just is being necessary in the Aristotelian sense. In this sense it is perfectly reasonable to require natural laws to be 'necessary'.

However, some time before Hume there began a trend to the modern notion of a necessarily true proposition as one the falsity of which is inconceivable. In this sense, laws of nature are not necessarily true—their falsity can be imagined. Thus it could have struck Hume as a discovery that causal connections are not logically necessary connections—a discovery rather than a triviality. This could have been all right. But Hume goes from saying that natural laws are not truths of reason to saying that they are not based on reasoning. He goes on to say that predictions about the future are founded, if at all, on habit. So the doctor who predicts that the patient cannot

[2] *An Enquiry Concerning Human Understanding* (The Liberal Arts Press, 1995), 46.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

recover without antibiotics ends up with the same generic kind of 'foundation' for his belief as the quack who says that the patient's only hope is a dose of his snake oil.

Of course Hume, being a profoundly intelligent and reasonable thinker, does try to mitigate the absurdity of his result. But he is swimming with an anchor—his assumption that all reasoning about predictions must have some basic principle that provides a common justification for all cases that are justified. He speaks of finding a combination to have held in the past (past A's have been B's) and then expecting future A's to be B's. He says this forming of expectation cannot be by any process of reasoning and that this 'principle' is not founded on reason, but habit. When we see lots of A's turn out to be B's, we develop the habit of expecting this connection to continue.

The observation that this is not true in general might not trouble Hume. He could say that some constant conjunctions lead us to form habits of expectation and some do not. But he has come to the conclusion that these expectations are all of them without foundation in reasoning. This is unwarranted, and may well be due to assuming that unless the principle of assuming future A's to turn out like past A's is a true principle about rational inference, no 'inference' fitting its form can be anything but a habit of expectation, rather than a rational inference.

Hume's pessimism about the prospect for finding a rational basis for expectations about the future or even about the present behaviour of physical objects is hard to assess because he does take some trouble to allow that some expectations are wiser than others. He does not endorse the view that the quack's habits of expectation are just as good as the distinguished physician's. In fact, when it comes to respecting reputation, he is quite conservative. He just avoids calling the wiser expectation one better founded in reasoning. This raises the possibility that he is engaging in an idiosyncratic use of 'reason' which would leave the translation of his views into more natural everyday language in the hands of scholars.

One prominent summarizer of Hume's epistemology is W. V. Quine. Quine characterizes Hume as having held the justification of conclusions about 'bodies' to depend on translating the claims in question into purely 'sensory terms' and then deducing them from other purely sensory claims which were not susceptible to doubt. Quine seems to regard Hume as having been somewhat successful in translating but a failure at making the required deductions.

> What then of the doctrinal side, the justification of our knowledge of truths about nature? Here Hume despaired. By his iden-

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

> tification of bodies with impressions he did succeed in construing some singular statements about bodies as indubitable truths, yes; as truths about impressions, directly known. But general statements, also singular statements about the future, gained no increment of certainty by being construed as about impressions.[3]

In addition to this view as to the nature of Hume's thinking about justification, Quine seems to express some considerable agreement with the ideas as he characterizes them:

> On the doctrinal side, I do not see that we are any farther along today than where Hume left us. The Humean predicament is the human predicament.

This seems to be saying that, among other things, Hume was correct to despair of justifying any singular statements about the future. Whether Hume did so despair is a scholarly question I will not pursue. There is, of course, the related scholarly question as to the proper interpretation of Quine's remarks. The quotations from Hume and from Quine conflict with what I take to be plain and obvious facts about justification and reasonableness. Without further enquiry as to whether their remarks really are meant to apply to these facts, I propose to consider an example of such a fact.

Suppose that a group of people have had an ongoing discussion of the prospects of Bill Boggs' winning the Boston Marathon. Most are sceptical of Bogg's chances, but Smith thinks he can win. Boggs was seventeenth two years ago and ninth last year and has been looking strong in recent competitions. Jones leads the faction which denies that Boggs has any real chance of winning. On the eve of the marathon, Smith is holding forth to a gathering of friends on his confidence that Boggs will win. Jones, arriving late, says grimly 'No, Boggs is not going to win tomorrow'. 'What makes you so sure?' says the cheerful Smith. 'Well,' Jones replies, 'Boggs was in a bad car crash this afternoon. His hip is smashed, compound fractures of both legs, a punctured lung—he is in a coma at the hospital in intensive care.'

Here we have a singular statement about the future—that Bill Boggs will not win the Boston Marathon tomorrow. It seems that Hume's position is that Smith's assertion of this has no basis in reason, and Quine's position is that it is correct to despair of justifying such a conclusion as Smith's. In so far as this is a true account of what these philosophers are saying, what they say flies in the face of such facts as that in this case, the singular prediction is entirely rea-

[3] 'Epistemology Naturalized' reprinted in *Naturalizing Epistemology* second edition, edited by Hilary Kornblith (MIT, 1994), 15–31, p. 17.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

sonable and fully justified. Jones is justified in being absolutely certain that Boggs will not win. He is at least as justified as he is in believing that all bachelors are unmarried or that 7+5=12.

This is not to say that Jones would be at all justified in regarding it as logically impossible that Boggs will win. If Jones is intelligent and well educated, he should be able easily to imagine a consistent scenario in which Boggs recovers overnight and goes on to win. Again, we can consider a second version of this case, in which Smith has had religious reasons for his claim that Boggs will win, claiming that God will help Boggs. If Smith responds to Jones' news by saying that he is convinced that God is going to heal Boggs tonight, then I, for one, would not question that it is in the power of God to produce such a miracle. Merely citing the severity of Boggs' injuries would not suffice to answer doubts based on such considerations.

Whether the fact that Smith has fervently prayed for the miraculous recovery of Boggs would provide a basis for doubting Jones' conclusion need not be considered further. In the case under discussion, the first one, there is no prayer for recovery or religious hopes for it. There is no doubt about Boggs not winning and there should be none. It is perfectly reasonable and justified to conclude that Boggs will not win. If asked for a reason for that conclusion, then the statement of Boggs' injuries is overwhelmingly adequate. One could question whether the reports are true. But to grant their truth and then doubt the conclusion would be irrational.

My use of 'fact,' here may well be challenged on the grounds that the case is purely imaginary. It is not intended to be purely imaginary, but primarily is a convenient version of countless actual cases which could be found. We could say that specific names have been changed to protect the innocent. The fact is that there are lots of cases relevantly like the one described. With respect to such cases, Hume's talk of the absence of a foundation in reason makes sense only to the extent that he means that the conclusions are not 'truths of reason', that is, necessary truths. And Quine's talk of 'despair', makes sense only to the extent that he means despairing of an absurd project. If the suggestion is that we should despair of attempting to justify such conclusions as that Boggs won't win, then I must despair of making sense of the suggestion.

However, we might move from the question whether the reasons offered for Jones, conclusion are conclusive, to the question, what reason we have for regarding them as conclusive. And then things can begin to look dark. Can we find a foundation in reason for holding that the reasons offered are conclusive? I do not know of any. But surely it is not just a brute fact that such reasons are conclusive?

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

I do not know. I claim only that they are conclusive, and that in the relevant cases, there is no basis for doubting that they are conclusive. No reason is needed for the reasons being conclusive. No reason could reasonably be requested.

It may be replied that even if it would not be reasonable to doubt the adequacy of Jones' reasons, it is always possible to change the subject, and ask, as a philosopher or general enquirer, why it is that reasons of the sort offered by Jones are so conclusive. You might explain that this is not intended at all to raise any doubt about the conclusion or the adequacy of the reasons, but just to express interest in the theoretical question as to the nature of the adequacy.

It seems quite reasonable to raise such a question, and it is no doubt a failing in that field of enquiry to answer as I do, that I do not know any answer. The fact that no answer is required as a condition of the justification of Jones' conclusion is not an answer to the further philosophical question. But it should none the less be a caution to those who do try the philosophical question, that it would be a severe defect in an answer if it had the consequence that Jones' conclusion was not justified after all. Unfortunately, it is common for philosophical discussion of justification and reasoning to ignore this distinction as to the limits of the enquiry.

It is even common to simply characterize 'induction,' as a kind of reasoning which cannot achieve certainty. One excellent and justly distinguished textbook may be cited merely because it is absolutely typical. It is said:

> A deductive argument is valid when its premises, if true, do provide conclusive grounds for the truth of its conclusion ... An inductive argument makes a very different claim: not that its premises give conclusive grounds for the truth of its conclusion but only that its premises provide *some* support for that conclusion.[4]

This could be a genuine distinction of kinds of arguments. Then Jones' argument would be deductively valid, and some sketchy proofs in number theory might be inductive. But it is also clear that the texts take it to be necessary and sufficient for deductive validity that it be logically impossible for premises to be true and conclusion false. On that meaning, Jones' argument is not deductively valid. Furthermore 'arguments' of the form 'P; therefore, P' are deductively valid even though the premise, even if true, provides *no* support for the conclusion. What is worst, though, is that the texts are

[4] *Introduction to Logic* by Irving M. Copi and Carl Cohen, ninth edition, (Macmillan, 1994), 57.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

drawn, by their general discussion, inevitably to the position, not usually acknowledged, that an argument such as Jones' *is* 'inductive', with the disastrous consequence that they are holding that in such an argument it is not claimed that the premises constitute conclusive support. It is further said of an 'inductive argument' that it 'is one whose conclusion is claimed to follow from its premises only with probability, this probability being a matter of degree and dependent upon what else may be the case.' (p. 60)

Of course it is true that in any argument, if it is claimed, as a premise of the argument, that the other premises provide conclusive support, and 'conclusive support,' is understood to mean in part that if those premises are true, then the conclusion is true, then the argument is deductively valid (which is not to say that the premises *do* provide conclusive support). That is, any argument of the form 'P1, P2, ... ,Pn, and if P1–Pn are true then Q is true; therefore Q', is trivially valid. And in any sincere argument, the arguer believes that if his premises are true, then the conclusion is true. If this belief is added as a premise, then the resulting argument is trivially valid. (If the claim of support is added as a *self-referential* premise, as in 'If this very claim is true, then Q' we get another matter, which is even less relevant here.) But this is not relevant here. I assume that what was intended was that for an inductive argument, a reviewer should not claim on behalf of the argument, that it was conclusive.

This is a bad mistake, which can only have a bad influence in teaching critical reasoning. A proper characterization of argument generically should make it a necessary condition of being a conclusion that reasons be offered for it as conclusive, that is, such as to justify drawing that conclusion. If the argument is not a good argument, then it may be false that the premises provide conclusive reasons. But it is essential to being an argument that the premises be put forward as providing conclusive reasons. To teach that an inductive conclusion is never really a conclusion should make 'inductive' a warning term, an indication of bogusness. But that is not what is generally meant by 'inductive'.

This is not, however, to express any confidence about what is generally meant by 'inductive'. For 'deductively valid' we have the impossibility of premises true, conclusion false. This definition has its critics. Some complain about the vagueness of 'impossibility' others about the concept's paradoxical insensitivity to relevance. Furthermore, it is a definition only of 'deductively *valid*', not of 'deductive'. It seems there are invalid deductive arguments.

By contrast, the definitions available for 'inductive', do not focus on the successful case. The most common definition (or small fam-

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

ily of definitions) is by reference to the practice of concluding that future A's will behave like past A's in respect of turning out to be B's. This is what Hume called a 'principle' and treated as needing, but not having, a basis in reason. Anthologized writers on induction generally relate themselves somehow to some version of this principle (PI). Here is Bertrand Russell's version of PI:

> (a) The greater the number of cases in which a thing of the sort A has been found associated with a thing of the sort B, the more probable it is (if no cases of failure of association are known) that A is always associated with B; (b) Under the same circumstances, a sufficient number of cases of the association of A with B will make it nearly certain that A is always associated with B, and will make this general law approach certainty without limit.[5]

Russell holds that this principle cannot be proved or disproved by experience, and that it is a fundamentally important principle.

> All arguments which, on the basis of experience, argue as to the future or the unexperienced parts of the past or present, assume the inductive principle; hence we can never use experience to prove the inductive principle without begging the question. Thus we must either accept the inductive principle on the ground of its intrinsic evidence, or forgo all justification of our expectations about the future. If the principle is unsound, we have no reason to expect the sun to rise tomorrow ... or to expect that if we throw ourselves off the roof we shall fall. (pp. 68–9)

This is an extravagant claim about the consequences of the principle's being unsound, considering that it appears on its face to be unsound. A famous example of Russell's own will suffice to bring out this appearance: 'The man who has fed the chicken every day throughout its life at last wrings its neck instead, showing that more refined views as to the uniformity of nature would have been useful to the chicken'. (p. 63) Useful or not, similar refinement is needed to save Russell's principle. Otherwise, we simply let A=a morning in the life of the chicken, and B=a morning when the farmer has fed the chicken and not wrung its neck. We need not enquire as to the thought processes of the chicken. The farmer and we philosophers as well all know that in a substantial number of cases 'a thing of the sort A has been found associated with a thing of the sort B'. Furthermore, 'no cases of failure of association are known'—remember that we are dealing with a specific case; that chickens generally do not fare well in similar arrangements has been left out in the formulation just quoted. It clearly follows from Russell's

[5] *The Problems of Philosophy* (Oxford, 1979), 67.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

principle that it has daily been getting more probable 'that A is always associated with B'. And this is absurd. At the very least, even if you refuse to admit that the farmer knew all along that he was going to kill and eat the chicken, (perhaps you appeal to scepticism about the future) you could not reasonably hold that the probability of this not happening, and the chicken's being fed indefinitely, is improving daily.

Russell does follow up with an argument which might be thought to answer this objection. He stresses that 'probability is always relative to certain data'. Thus a man who concludes in accordance with Russell's principle, after seeing a great many white swans, that it is *probable* that all swans are white, is not shown unreasonable by the discovery of black swans in Australia.

> A man might know that colour is a very variable characteristic in many species of animals, and that therefore, an induction as to colour is peculiarly liable to error. But this knowledge would be a fresh datum, by no means proving that the probability relatively to our previous data had been wrongly estimated.

Applied to the chicken case, the farmer would have additional data beyond the facts about regularity of feeding and non-wringing. The chicken might be the proper thinker for the example after all. It presumably has been given no reason to suspect that the farmer is anything but concerned to keep it well fed.

This defence of the chicken's reasoning is not compatible with the claim that 'more refined views as to the uniformity of nature would have been useful to the chicken'. The chicken was presumably doing as well as could be done on the information available to it. What it needed was more information about farmers. However, just how such information might have been used is left completely undetermined by Russell's principle.

Russell's defence is unclear, but the best that can be made of it is to take it as restricting the application of his principle to cases in which you have no other relevant information than that past A's have been B's. This gets him out of the farmer case, because the farmer has other relevant information. But then Russell's claim that the reasonableness of our expectation that the sun will rise tomorrow depends on his inductive principle is indefensible, since it is clearly false that the only relevant information we have about that is that it has always happened in the past. Just as we can use our relevant background information to predict that the chicken will not get fed today despite the fact that it always has, we (rather, astronomers) can use relevant background information to predict that the sun will rise tomorrow.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

Furthermore, the picture of someone knowing only that past A's have been B's is deeply obscure, so that the principle not only does not apply in everyday cases; it is not clear that it applies in any cases. We may try to imagine ourselves in a situation like the chicken's, prisoners of a power we do not understand, which brings us food every day. Is it reasonable for us to infer that the power will continue to do so? We may as well go out to the feeding place when hungry—what else is there to do? If escape were an option, it would be a different matter. We may have no basis whatever for theorizing about our food supply.

Russell's principle is either trivially false or useless. It is trivially false taken literally. Modified so that it licenses concluding the next A will be a B when past A's have been, etc, and the other available background information is favourable, it is useless. Such vague qualifications as 'relative to proper background information' or 'given an adequately large sample' can be interpreted so loosely as to make the 'principle' trivially immune to counterexample, but question-begging to apply and hence useless.[6] To criticize Russell's position adequately requires giving due weight to this fact. But it is difficult to get this weight right.

Hans Reichenbach attempted to defend the PI by appeal to a frequency definition of probability.[7] The probability of an A being a B is the limit of the ratio of past A's that are B's to past A's, as the number of A's observed goes to infinity. Our best posit, according to Reichenbach, is to take the ratio observed so far as the probability. He stresses that we won't know this is the true limit ratio; but *if* there is such a limit, this is our best hope of getting it right.

This is false. To base my estimate of the probability of heads with a given penny on the observed ratio after, say, 10 tosses, is to be likely to be mistaken about the true probability. It may be replied that we know this only because we know of trials that were sufficiently long which established the probability. But that is not how we know the probability, or not the only way. One omission in the PI is any specification of a sufficient number of trials for setting a given probability level. Without that the PI is useless. It is useless, but in making this clear it helps to note that the probability that all baboons of a certain species have brightly coloured buttocks, based on seeing four healthy specimens, an adult male, an adult female, and a young female and an old male, can be set close to 1 with more justification than for setting the probability of heads based only on

[6] I have discussed the trivializing effect of qualifications to the PI in 'On Having Reasons', *Analysis* 1966, 189–92.

[7] Page references are to the excerpt from *Experience and Prediction* which is included in *The Theory of Knowledge* edited by Louis Pojman.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

10 tosses. We do of course make judgements about how big a sample should be. But we do not have a general principle covering all such judgments.

The idea that the limit at infinity will be reached by Reichenbach's approach if reached at all is cold comfort indeed, when we consider (1) it *won't* be reached at all, being infinite: and (2) *any* finite sequence of outcomes, however long, is mathematically compatible with *any* limit probability.

Reichenbach's characterization of 'the aim of induction' is to find a probability in the sense of the mathematical version of the frequency theory—a limiting ratio. His thought is that if you always assume that the probability of an A being a B is the ratio of past AB's to past A's, then if there is a limiting ratio, you will find it this way. This is a mathematical truth. But you will not get to infinity, and there may not be a limiting ratio waiting if you could. And the fact that the probability of an A being a B does turn out to be n (when the limit exists) may be a pitifully poor guide to deciding whether *this* A is going to be a B, since this A is also a C, D, etc. And if this isn't all bad enough, Reichenbach tells us that his approach is all the hope we have for rationality in the prediction business. Even habit looks good by comparison.

In a well-known paper[8] Paul Edwards criticizes Russell for holding 'that unless we appeal to a non-empirical principle which he calls "the principle of induction"' we must answer in the negative such a question as

> (1) Assuming that we possess n positive instances of a phenomenon, observed in extensively varied circumstances, and that we have not observed a single negative instance (where n is a large number), have we any reason to suppose that the n+1st instance will also be positive?

Edwards argues 'that question (1) can be answered in the affirmative without in any way appealing to a non-empirical principle'. To this end he cites the example of inferring that someone who has just jumped from a window on the fiftieth floor of the Empire State Building will fall towards the street rather than upwards. This example should be compared with the one given earlier about a severely injured man not winning a marathon. That one was an overwhelmingly reasonable inference, one not open to any reasonable doubt. The only problem with the jumper example lies in characterizing it as an inference at all. People have been observed to jump from the Empire State Building. But it is not obvious that

[8] 'Bertrand Russell's Doubts About Induction', *Mind* LVIII (April, 1949), 141–47.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

horrified onlookers 'inferred' the jumper would fall. Rather, they saw that the man was going to fall. (If the man had been wearing a parachute, and there was a windstorm going on with a tremendous updraft, then the case would be different.) It is obvious that in some cases of Empire State jumpers, expecting a fall is right, and actual cases can be cited. That this would be inferential is not so clear, though cases can be imagined in which this too would be the correct description. Recall Russell's claim: 'If the principle is unsound, we have no reason to expect that ... if we throw ourselves off the roof we shall fall'. This, as was argued, is false. The principle is unsound (if not hedged so as to be useless), and we do in some cases none the less have reason to expect, etc.

This does allow a positive answer to (1) based on empirical observations. Suppose we are observing a man on a ledge on the fiftieth floor, about to jump. We say 'Oh no!—this will be the 17th case of someone jumping—he's going to fall to the street just like the previous 16!' This is an empirical claim and is sufficient to warrant an affirmative answer to (1) *interpreted as prefaced with* '*Does it ever happen that*?'. For this is a case in which we have observed n(=16) positive instances and no negative, and we have reason to suppose that the nth(=17th) will be positive.

It might be objected that in this case we do not have 'extensively varied circumstances' since all our previous cases were from that same building. This is worth noting in connection with the vagueness of such hedging phrases, but we can override the objection for present purposes. For we could add all the known cases of falls from buildings (while allowing for ascending parachutists and the like) to qualify as having positive instances observed in extensively varied circumstances. This should be enough to establish Edwards' claim that 'without in any way calling upon a non-empirical principle for assistance, we often have a reason for supposing that a generalization will be confirmed in the future as it has been confirmed in the past'.

However, this is assuming that the ambiguous quantification of an affirmative answer to (1) is taken as the modest 'It does happen in some cases that, with n positive instances, etc. we have reason to expect n+1 to be positive'. It is very different if we read the affirmative answer as the claim that *whenever* we have n instances in extensively varied circumstances, etc. then we have reason to expect n+1 to be positive. We may call the former, existential reading of an affirmative answer to (1), E1, and this latter, universal reading, U1. U1 is just another version of PI. To argue that U1 is an empirical truth is no improvement over Russell's claim that PI is an indispensable assumption. The question whether PI is 'empirical,' is of

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

course another topic for philosophical disputation. But the dispute would seem less urgent if it is recognized that PI is not even a true generalization.

It seems clear that Edwards does not intend to argue for PI as a general truth. He speaks of showing that 'we often have a reason' never of always having a reason, when we have the n positives, etc. His examples and claims are eminently sensible. But he is not successful in extricating inductive justification from dependence on PI. He centres his argument against Russell on the idea that Russell has tacitly given a high redefinition of 'reason for an inductive conclusion', to mean 'logically conclusive reason'. And he replies that this is not what is meant, and offers a characterization of what is the 'main sense' of 'reason'.

The charge against Russell is incorrect. Russell does not deny that we have reason in cases where we lack logically conclusive reason. Rather, he says that our having reason in such cases is dependent on PI, and he says that PI is not capable of being proved or disproved by appeal to experience. Edwards picks Russell up on his claim that PI is nonempirical. This claim is hard to evaluate, and I have not attempted to do so. My objection was focused on the dependence claim. Being false or trivial, PI is useless in justifying inductive inferences, so that none depend on it. Edwards slips (unintentionally, I believe) into defending an affirmative answer to (1), as if it were analytic. That makes his modest preference for E1 over U1 indefensible. About the meaning of 'reason for an inductive conclusion,' he says:

> According to this main sense, what we mean when we claim that we have reason for a prediction is that the past observations of this phenomenon or of analogical phenomena are of a certain kind: they are exclusively or predominantly positive, the number of positive instances is at least very large, and they come from extensively varied sets of circumstances. This is of course a very crude formulation. But for purposes of this article, it is, I think, sufficient. (p. 358)

Now, if that is the definition of the relevant meaning of 'reason', then an affirmative answer to (1) comes to this:

> Assuming that we possess n positive instances of a phenomenon, observed in extremely varied circumstances, and that we have not observed a single negative instance (where n is a large number), then the past observations of this phenomenon or of analogical phenomena are exclusively or predominantly positive, the number of positive instances is at least very large, and they come from extensively varied sets of circumstances.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

To say that this is 'often' true, or true in many cases, is absurd. On this interpretation, (1) is properly read as U1 and E1 is confused. Note that 'E1' and 'U1' are only abbreviations for disambiguation of the quantification of (1). 'Reason' was left uninterpreted. Edwards' account of the 'main sense' of 'reason' is what makes U1 express a truth and E1 an oddity (rather like 'Some contradictions are false'—of course the trouble is not that E1 is *false* on this interpretation).

Edwards is highly sensitive to the variety in use of 'reason'. He notes, for one example, that someone might use 'reason' merely to apply to a positive instance of a generalization. This suggests that he is concerned with something more, in his 'main sense ' : justifying reason. The topic is justifying inductive inferences, not showing that we can say that induction sometimes provides some, even if inadequate, reason. But even in the form of establishing just 'some reason', Edwards' claim about meaning is either false or trivial in the way PI is.It is the claim that the mere fact that past A's have been B's, with no further information, constitutes adequate reason to justify concluding that the next A is more likely than not to be a B. My view is that this is never true, if we do not consider the background information we have associated with our conception of A's and B's. And even considering such information, it is only sometimes true that past instances alone, under only one concept A, give reason to expect the next A to be a B.[9]

I agree with Edwards that there are lots of cases in which, having n positive instances, etc., we are justified in expecting n+1 to be positive. I believe that Russell did not deny this either. But contrary to both Russell and Edwards, I hold that in such cases, the reasons are not the fact that there have been n positives, etc. For example, when Jones justifiably predicts that Boggs will not win the Boston Marathon, his reason is that Boggs has been broken up in a car crash. That is what justifies his conclusion. Its not that n people have been observed broken up in that way and have not functioned well shortly thereafter, and n is large and the cases are varied, so we may conclude there is reason in this, the n+1st case. That might be undertaken by someone trying to prove that Jones' reason is a good reason, which is a risky exercise perhaps quite honourably undertaken by philosophers, though often with bad results.

It would seem that asking about the meaning of 'reason,' or 'inductive reason,' would help in clearing up puzzlement about induction. But if I am right, it is not easy to do this successfully, as the case of Edwards should suggest. Another philosopher who

[9] For further argument on this point, see 'On Having Reasons'.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

approaches 'the problem of induction' in terms of meaning is P. F. Strawson, who says: 'The only way in which a sense might be given to the question, whether induction is in general a justified or justifiable procedure, is a trivial one which we have already noticed. We might interpret it to mean 'Are all conclusions, arrived at inductively, justified?' i.e., Do people always have adequate evidence for the conclusions they draw?' The answer to this question is easy, but uninteresting; it is that sometimes people have adequate evidence, and sometimes they do not.'[10]

'Sometimes they do and sometimes they don't'—this suggests an answer to Hume's question 'What is the foundation of all conclusions from experience?'(434), namely 'Some conclusions from experience do not have foundations, being unfounded mistakes—and as for the "basis" of those that have bases, *it depends on the particular case*'.

However, it is all too easy to stray from this path of particularism back into the puzzles of generality. Strawson's use of 'i.e.' in the above quote suggests his rewording of (1) 'Are all conclusions arrived at inductively, justified?' should be (2) 'Do people always have adequate evidence for the *inductive* conclusions they draw?', to which his answer would be 'Sometimes they do and sometimes they don't'. But he immediately continues '"forming rational opinions about the unobserved on the evidence available" and "assessing the evidence by inductive standards" are phrases which describe the same thing' (p. 258). If this were right, or were a stipulation as to how Strawson will use 'inductive' then the answer to 2 would have to be 'yes, trivially, by the definition of the terms'. That is at odds with 'Sometimes yes, sometimes no'. Strawson might say that he is talking about two different uses of the term 'inductive'; one meaning 'rational', the other meaning, say, 'concerned with the unobserved'. If so, he did not make this sufficiently clear, and there is evidence here of a temptation to use the same term in conflicting ways.

Besides the equation between being a rational method for assessing the evidence and being an inductive method, Strawson also emphasizes an important relation between being a *successful* method and being an inductive one. He speaks, in reference to a method, of asking whether its employment is inductively justified, whether it commonly gives correct results' (p.258) as if these are the same. Again, 'by the very fact of its success it would be an inductively supported method'. Finally, he says:

> So every successful method or recipe for finding about the unobserved must be one which has inductive support; for to say that a

[10] *Introduction to Logical Theory* (Methuen, 1952), 257.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

> recipe is successful is to say that it has been repeatedly applied with success; and repeated successful application of a recipe constitutes just what we mean by inductive evidence in its favour ... any successful method of finding out about the unobserved is necessarily justified by induction. This is an analytic proposition. (p. 259)

In addition to this analytic truth about being inductive, there is also said to be an analytic connection between being an application of the 'straight rule' and being inductive. Strawson says

> ... it is an analytic proposition ... that other things being equal, the evidence for a generalization is strong in proportion as the number of favourable instances, and the variety of instances in which they have been found, is great. So to ask whether it is reasonable to place reliance on inductive procedures is like asking whether it is reasonable to proportion the degree of one's convictions to the strength of the evidence. Doing this is what 'being reasonable' *means* in such a context. (pp. 256–7)

I will assume that Strawson means that *known* favourable instances support a generalization. If S has no good reason to think past A's have been B's, then the fact that they have been need not guarantee that it is reasonable of S to believe future A's will be B's. But Strawson does appear to hold that the more you know past favourable instances, etc. then the more justified confidence you may have that the future instances will go the same, and that moving to such confidence levels is proceeding inductively.

This suggests three equations. In formulating them I will use 'method', as short for 'method or recipe for finding out about the unobserved'. We have

E1. being a rational method=being an inductive method
E2. being a successful method=being an inductively justified method
E3. being an application of the straight rule=being inductive.

All of these are represented as being analytic. There is some reason to take Strawson as intending E2 as only a one way entailment from 'successful' to 'inductive'. Anyway, I will only use it that direction.

Strawson calls each of his three equations (or two equations and an implication) 'analytic', that is, truths about the meanings of the terms involved. This could help give being analytic a bad name. It is interesting to note that each of these three equations could be said to merely reflect facts about how the terms are used. But then the terms are seriously *ambiguous*. This is a danger in the use of the 'analytic method'. Note that the equations amount to a strong

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

defence of the PI. Following the PI is by definition being inductive and also is necessarily rational and is at least a necessary condition of being successful at forming opinions about the unobserved on the basis of past observations. (What then, has become of Strawson's earlier warning that the only sense available for the general question whether inductive inferences are justified is one which makes the appropriate answer 'Sometimes they are and sometimes they aren't'?)

Strawson gives an example in which someone uses a wild 'method'—closing his eyes and guessing and 'it's usually (always) the right answer' (p. 258). But in his case, it appears that the guesser has observed over a long term that this method does work for him. It is hard to imagine this really happening for a completely open range of questions. The guesser would be finding himself to have godlike powers. But suppose we confine it to guessing whether a certain coin will land heads. As a 'method' for forming opinions about how the coin will land, guessing, at the outset (in a normal case with no special antecedent information) would not be rational. Guessing would be as good as any other method, but not good. If forced to bet on the coin, then guessing would be a reasonable thing to do, but being sure you would be right, and being willing to act accordingly, would not be rational.

If you were forced to bet, or just idly tossing the coin, and you noticed that your guesses were always right, it would be reasonable to find this quite surprising. You might privately test the coin and the ridiculous hypothesis that your guess is reliable. You might 'find' that your guess is reliable for this coin and only this one. That is, you might find that trial after trial your guess is right, and that this does not happen for other coins. You could get in the position of having good reason to believe the extraordinary generalization suggested by such observations. You would also have good reason to be very cautious about talking to anyone about it, knowing that they would reasonably regard you as nutty.

If you then were required to predict the outcome of a series of tosses of your coin, and you got the success rate you would reasonably expect and no longer find surprising, an amazed observer might ask you how you were doing this, what method you were using. You could then reply truly, 'My method is just to close my eyes and guess'. It might then seem that it must follow that this method is rational and also successful and thus a confirming instance of one of Strawson's equations.

In my opinion, this would be a mistake. Your method for arriving at a prediction should not be identified with your method for arriving at a rational opinion about how the coin will land. The idea of a

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

method for acquiring opinions makes some sense, but it is hard to make sense of guessing as being such a method and it is surely not a method of justification. Your justification is not merely that you so guessed, but all the considerations that have led you to believe that your guesses about this coin somehow turn out correct.

It might be replied that, none the less, the method of guessing is a rational one to use, and a successful one, so the equation in this case is not in question. This would overlook the point that it is not a method for justifying or founding opinions about the unobserved, or even a method for forming opinions. Guessing does not require believing (it can of course include believing, but that is in this case merely a distraction I will not bother to work through). It is in this case just a method for arriving at predictions, and predictions should not be confused with judgments.

It must be admitted that this distinction is especially easy to lose track of in connection with predictions to oneself, since philosophers are sometimes tempted to regard believing as saying to oneself. This makes it difficult to imagine how someone could privately experiment with the coin and learn that his predictions are always correct, and only then come to *believe* they are always correct. For his sayings to himself, that is, his predictings, would be confused with beliefs by the inner saying model. We could, of course, have it that the man is actually forming opinions about how the coin will land. Since we are dealing in mere logical possibilities, it costs nothing extra to have it that way. But it is better to stick to what is minimally required to be guessing. One can say to oneself 'Its going to land heads'. An audience helps make this into a prediction, but is not essential.

To give an analysis of 'A guessed the answer to the question whether P' would be a difficult project. Suffice it to say that in some cases, one guesses merely by forcing himself to answer the question when he has no opinion as to the truth of the matter. The coin is covered, and you are asked whether it is heads and you say 'Okay, I'll say it isn't, this time'. It is not surprising that you may turn out to be right. But if you turn out to be right for an extraordinary run, that is extraordinary. Then, knowing that that has happened may give you reason to think that somehow, your guess is a good indication of the state of the coin. Now, the fact that you said 'It's heads' is good evidence that it is indeed heads. But it is incorrect to say that so saying is a method for forming a belief on the question whether it was heads. You form the belief by considering, among other things, the fact that your guessings have been extraordinarily accurate.

We may say, then, that by the method of guessing, someone has

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

been successful at making true predictions, but this is not a license for saying that guessing has been a basis for forming justified beliefs about what will happen. It is knowing about the success of the predictions and having reason to doubt that the past success can be explained as mere coincidence that justifies believing that the predictions will continue to be true. And that is not mere guessing.

What if our guesser is completely convinced he is getting it right, but without observing that he is? What is observing here? He tosses the coin and doesn't even bother to look at the result. He is so sure his guess is right that it seems to him pointless to bother to check (and he thought it pointless from the outset, not just after developing reasonable confidence based on his success rate). Like the man who has checked and found the coin to be landing as he guesses it will, this one believes he has had a long run of successes. The difference is that this belief is not reasonable, so that the conclusion based on it, that he will continue to be right, is not reasonable either. The guesser is not learning of his success, just having it. Surely this does not qualify his procedure as rational. If it qualifies as following the PI, then so much the worse for the PI. If someone forms the unreasonable opinion that past A's have been B's, is it then reasonable of him to infer that the next A is likely to be a B? Whichever way this is answered, the case is a problem for Strawson's equations. For even if to qualify as *following* the PI requires knowing how things have been turning out in past cases, this man is just as successful as if he were being rational and following the PI, so that E2 entails that he is following the PI. Ruling out the case would thus only involve further inconsistency.

This shows that the three equations cannot be correct. Our lucky guesser is following a 'method' for making judgments about the unobserved on the basis of his alleged 'observations' which is successful and thus inductive by E2. Whether or not it qualifies as following the straight rule it clearly should not qualify as rational. Strawson's three equations cannot all be truths about one and the same meaning of 'inductive'.

This criticism was based on just the one way reading of E2. It is interesting, in passing, to enquire whether, in addition to its being possible to be successful without being rational, it is possible to be rational without being successful. However, this question is clouded by the Kantian point that success in forming expectations about the future is a necessary condition of being a normal person. To recognize your surroundings, on this view, requires expecting things to behave as things of those kinds should. Just walking down the street, you expect that the pavement will not rise up and seize you, or the ground give way under your step, or your foot turn into a

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

dragon, etc., etc. If these 'inductive expectations' were not 'successful' you would not exist as an integrated personality.

There is a lot of truth in this picture. It provides some basis for the editorial quotation which began this paper ('Our existence ... is based on the principle of induction'), but it overintellectualizes walking down the street. It is perhaps true that one expects, e.g., that his head is not going to turn into a pumpkin, and even true that there is some justification for this expectation. Still, if someone is asked to justify such an expectation, he will be at a loss. And he should be. The philosopher may be able to work up a patter purporting to rationalize the expectation, and the psychologist may produce a correct causal account of how normal people come to have such expectations, but the fact is that such expectations are in the overwhelming run of the time not occasions for reason-giving at all and are not cases of forming expectations on the basis of reasons. That countless expectations are not shattered (as they are when one meets up with a mugger, etc.) is indeed a condition of our healthy existence, but we should not characterize this as the requirement that induction be justified or even successful.

Strawson says that 'to say that a recipe is successful is to say that it has been repeatedly applied with success'. Then no recipe could be successful on its first application. Presumably it would not be unsuccessful on the first application either, no matter what happened. So what would it be for a recipe or method to be unsuccessful? It would seem to be being repeatedly applied without success. The argument runs: 'repeated successful application of a method constitutes just what we mean by inductive evidence in its favour' and repeated successful application is also what we mean in saying the method is successful; therefore a method which qualifies as inductive will be successful.

This argument has already been challenged on the grounds that the method of relying on a method which you know has been repeatedly applied successfully is not automatically the same method as the one you have been repeatedly applying successfully. The method of relying on long term successful methods may be made out as the method of following the straight rule, where in this case the past A's that have been B's are past applications of method X which have been successful. It was observed that method X may well not be based on the straight rule itself and may be a method it would be idiotic to trust independently of considerable evidence of its successful application. The successful long term application of a method X does not prove that method X, in itself, is a rational one to employ. It is just that it may in certain possible cases be rational to employ methods, however silly they may seem, if they are known

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

to work over a long run. But appeal to that need not be appeal to method X.

Now we need to turn to the question as to what we can conclude from repeated unsuccessful application of a method. Could we conclude that the method is not a rational one to employ? Suppose that someone applies the straight rule to conclude the next A1 will be a B1. To qualify as employing this rule he has to have observed a sufficient sample of A1's and found them to be B1's. He has not just jumped to expecting any old A1 to be a B1 without careful collection of 'enough' examples. But despite all his excellence in sampling, he is mistaken—the next A1 is not a B1. This does not show he was not rational, nor that the straight rule is unsuccessful.

He then turns to the question as to whether A2's are B2's. Again he does a painstaking job of sampling and finally concludes the next A2 will also be a B2. Wrong again! But who's to blame him? This keeps up until one day, having seen an impressive number of A78's turn out to be B78's, and preparing to venture a prediction, he thinks 'Wait a minute! According to my records, this is the 78th time I have applied the straight rule, and in the previous 77 cases, I got beat. This darned rule just isn't working for me.'

What can we say to this sensible but unlucky soul? Suppose that we had been checking each one of his (inductive?) inferences to see whether his sample looked good. We have always found he was doing a terrific job, but have not stayed around to see how it turned out. It would be a job to make this story plausible, but maybe we have lots of predictors to check up on and we have been ordered just to check the basis and then get on to the next customer, etc., etc. We would then be quite reasonable in expecting that most of this man's predictions have been right. We might in fact be stunned to hear they were *all* wrong, perhaps incredulous. Should we conclude that he has not been proceeding inductively, or that his reasoning has been bad, or that he has been applying a bad method?

We might try the observation that our inductor has had countless successes, going back to the idea that his expectation that trees will not jump out of the ground and dance on their roots this time is based on the past string of cases in which they have not behaved that way, etc., etc.. This should be ruled out by our emphasis on cases in which an argument is actually presented, accompanied by the explicit assertion of a conclusion.

We could complain that, knowing he has failed in 77 tries, he should have begun sooner to factor in that consideration in the arguments he has been presenting us for review. We might not have endorsed his 70th prediction so enthusiastically if we had known

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

how the previous 69 had turned out. One might even try to avoid granting that he has been following the straight rule and being inductive, on the grounds that he should have been keeping track all along of how it is going with C's being D's where C=application by me of the straight rule and D=successful application.

This reply involves a shift in the interpretation of the PI, to include taking into account the problem of the reference class. In considering whether X is a B, you may not simply note that X is an A and survey past A's. You must also note whether it is a C and consider the performance of past C's. You might as well also be expected to decide which reference class, A's or C's, is the best guide to X. On this interpretation, it is indeed impossible for the PI to be applied unsuccessfully. But this is arranged by trivializing the idea of 'application of PI'.

I will defer to this so far as to refrain from claiming that it is possible for the straight rule to be unsuccessful. But this is only due to the obscurity of the rule. Once we bring in the 'problem of the reference class' it is a misleading pretence to speak of following the straight rule. When it is just a matter of past A's being B's, then it is possible that someone following the rule could have a long run of failures, once we stress that the applications are not in forming the expectations that are the source of the unity of experience but rather, formal pronouncements of predictions. But if we require that in asking whether this A, x, is going to be a B, we consider, not only the past record of A's but of any other class to which x belongs and which may be relevant, then the procedure is altogether different. We make it possible to rule out the possibility of long run failure by building an escape clause in the procedure for determining what counts as an application in the first place.

The 'problem of justifying induction' is sometimes characterized as the problem of justifying believing that the future will be like the past. Whether all induction is concerned with the future is a cloudy question we do not need to answer. It may be argued that even when we inductively conclude, say that primitive people inhabited a certain region, this essentially involves some prediction about how future observations will turn out. This seems doubtful, but need not detain us. It is enough to observe that the slogan that the future will be like the past (call that FPI) is very like the PI in the role it is mistakenly thought to play in inferences about what is going to happen. Paul Edwards argues as follows:

> In the ordinary sense of the word 'future' therefore, what Russell calls past futures *are* futures. They are futures in relation to certain other periods which preceded them. Now, the appeal to the fact that past futures resembled past pasts and past presents con-

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

> stitutes excellent inductive evidence for the conclusion that the future will resemble the past and the present. (p. 359)

This is, like Edwards' answer to his question (1), an answer that turns out, however intended, to be a generic defence of the principle FPI, and thus an unindiscriminating one. Edwards would certainly point out that the future often differs in striking respects from the past and that not all expectations of similarity are justified. But his argument, aimed at Russell, applies to just the sort of generic justification he elsewhere avoids claiming. On the question of FPI, he endorses the argument of an earlier paper by F. L. Will. Will's primary concern is with scepticism about the future. Future A's are likely to be B's because past A's have been B's, and the future will be like the past—in that respect at least. Hume often speaks in these terms, worrying 'that the past may be no rule for the future' (435). I have been arguing that it is a mistake to suppose that the rationality of predictions depends on the truth of such a general rule as PI or FPI. Will argues that it is a mistake to think that 'the future is forever hidden behind a veil'. He would say that 'the future is constantly being revealed', contrary to Russell's claim that 'We have experience of past futures, but not of future futures, and the question is: Will future futures resemble past futures?' (440)

Will's disagreement with Russell is apt to get lost in this question of whether we ever actually encounter the future in its full futurity. Better to consider some actual question about the future, such as whether Boggs will win. We all know he won't. Is the fact that the event is future a defeater to this knowledge claim? Is the reply 'It hasn't happened yet' a rational response, or just as rational as the claim that he won't win? No. Whatever Russell's position was, any position which entails that 'It hasn't happened yet' would suffice to show that one expectation was no worse than its contrary, both being nothing better than habits of expectation with no basis in reason, is just silly.

Will compares Russell's position to that of someone who makes a promise about next year and then a year later, in reply to the complaint that the promise isn't fulfilled, says 'This isn't next year'. But Russell has not arranged it so that the future never comes. It's just that any future year, from the perspective of a predictor, is different from any he has observed, in not being observed (yet). Will says: 'The correct conclusion to be drawn from the fact that time passes is that the future is constantly being revealed and that, in consequence, we have had and shall have the opportunity to learn more and more accurately what the laws of nature's behaviour are and how therefore the future will be like the past'. (449) Russell does not deny that the future is constantly being revealed. He only denies

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

that what we see of it could provide a basis for knowledge about the unrevealed part, unless we can appeal to PI.

One possible distraction is that Russell, when writing about the metaphysics of time, denies that 'past' 'present' and 'future' mark genuine features of reality. This need not concern us here. He will grant that 'we' exist at a certain time and that relative to that time there is a 'future', etc. Russell's rejection of an absolute future is, if relevant at all, just reason to say that he would grant Edwards and Will that there is no intrinsic difference between past futures and future futures. There is some difference between what has been observed and what has not yet been observed. Sometimes a sample of observed things is a good guide to how things are with respect to some as yet unobserved things. Russell does not deny this any more than Will does. The question is as to what role the PI plays in this guidance.

Russell's claim that the rationality of belief in general laws is 'completely dependent on the inductive principle' is false. The 'inductive principle' is an overgeneralization of no use by itself in forming expectations and laws. But Will's reply is unfortunately directed (or at least, not at all clearly *not* directed) at establishing that the inductive principle is, after all, a good one. Will is right in saying that 'we have the opportunity to learn how the future will be like the past,' and even to know this in advance. But this is not achieved or even facilitated by assuming that 'nature is uniform' or that 'the future will be like the past'. In some ways it will, in some it won't. 'The inductive principle' does not help in determining which.

Will is sensitive to the point that the general principle is dubious. Of the closely related 'Principle of the Uniformity of Nature', he says: 'It is ... difficult to interpret this so-called Principle in such a way that it makes a statement which is both definite and is not at the same time refuted in some areas of experience'. (444) This is an understatement. PI is simply a false generalization, so that searching for its basis in reason or deploring its lack of one is deeply misguided. This can be obscured by confusing the question whether the principle is true in general with the question whether any particular cases of it are reasonable. It is true that it is often reasonable to regard the next A to be likely or sure to be a B on the grounds that all past A's have been observed to be B's. The rationality of a particular inference following the form of the principle is not guaranteed merely by its being of that form, but other background information which has not been represented formally.

The mistake of thinking that all instances of PI are justified merely by being of that form leads to a further mistake. One becomes aware of the first mistake and moves to the view that *no*

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

instances of PI are justified. The idea would be that *if* 'induction' is to be justified it must be by the generic power of PI. Once it is seen that PI does not have the required generic power, it is then concluded that nothing of that form is justified unless some other generic source of authority can be found.

One boost to this general scepticism is the following fact: any time anyone appeals to the fact that all past A's have been B's to justify concluding that all are or the next will be, there is a contrary conclusion that is equally supported by the same facts! Let a B* be defined as a thing that is either examined before t and a B or not examined before t and not a B. Now, whenever all observed A's are B's, for the right choice of t, it is equally true that all observed A's are B*'s. So PI equally justifies the conclusions, both that all A's are B's, and that all A's are B*'s. Similarly for the next A being a B vs. its being a B*. But these conclusions are incompatible, being logical contraries. For the next A to be a B*, since it is now after t, it must *not* be a B.

This is a logically condensed case of the general problem of the reference class. You want to know whether X is a B. You know its an A, and they have all been B's. But its also a C, a D, etc. and they each have different ratios of being B's. So which one should you choose as your guide? Should you just average them all? A better suggestion, in my opinion, is to note that the description of the background considerations that make one reference relevant and another not are resistant to characterization with formal generality. This 'context dependence' view goes back to Aristotle and is behind the general scepticism of rationalists about the applicability of universal exceptionless rules to concrete particular cases.

Goodman's famous 'grue' is just one case of a B*. Goodman is a distinguished critic of PI. But he shares with the philosophers discussed earlier a tendency to overgeneralize in assessing the significance of the principle. Goodman describes a 'new riddle of induction' as 'the problem of distinguishing between projectible and non-projectible hypotheses'[11]. This projecting is just inferring and a projectible hypothesis will either constitute, or allow easy construction of, a justified inference about the unobserved, what would have been called a justified inductive inference. So Goodman, after making very clear that the PI will not do, is back to seeking a generic justification for induction.

His answer appears to be as follows:

> To speak very loosely, I might say that in answer to the question what distinguishes those recurrent features of experience which

[11] Nelson Goodman *Fact Fiction and Forecast*, 2nd edition, (Bobbs-Merill, 1995), p. 83.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

> underlie valid projections from those which do not, I am suggesting that the former are those features for which we have adopted predicates that we have habitually projected. (p. 97)

Goodman bases this conclusion on his famous example of a B*, the predicate 'grue', meaning 'observed before t and green or not observed before t and blue'. Before t, all emeralds have been both green and grue, so that PI gives the basis for incompatible projections about future emeralds. Goodman is suggesting that the only thing that justifies us in predicting the next emerald will be green rather than grue (and thus, for the right choice of t, not green, but blue) is that we are in the habit of using the term 'green' in predictions, while we are not in the habit of using 'grue'. That is, he goes back to Humeanism.

Qualifying the PI so as to project only predicates we have been in the habit of projecting successfully in the past would reduce the number of unreasonable projections that would be sanctioned. However, the fact that a predicate is one we are in the habit of successfully projecting does not protect us against the risk of making a foolish mistake about projecting it in some case. And the qualified PI would be just as mistaken as ever about the source of the rationality of such projections of ours as are rational ones. Scientists may coin a new predicate expressing a never before noticed property, say, of atomic particles. They may then predict that all particles of Kind K will satisfy this predicate, say, 'alpha'. To project instead a B* for 'alpha' could be seen to be irrational with no reference whatever to our past habits with regard to the projection of 'alpha', there being no such habits.

Against Goodman, it may be objected that 'grue' is not a purely qualitative property. For past values of t, emeralds have turned out to be green and not grue after t, the arbitrary cut-off date built into 'grue'. Goodman replies 'I simply do not know how to tell whether a predicate is qualitative or positional except perhaps by completely begging the question at issue'. He gives an argument: 'grue' and 'bleen' are indeed defined in terms of a cut-off date. But 'green' and 'blue' can be defined in the same way. For example, X is green iff it is either examined before t and grue, or not examined before t and bleen. Goodman concludes 'Thus qualitativeness is an entirely relative matter and does not by itself establish any dichotomy of predicates.' (p. 80)

This is a bad argument.[12] The fact that it is possible to define a

[12] I have criticized this argument in 'On Goodman's Riddle of Induction' *Ratio* 1970, 144–8.

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

predicate by reference to a cut-off date does not establish that determining the date is necessary to verify that the predicate applies. It is this latter feature that establishes that a predicate is not purely qualitative. It is possible to tell that something is green even if you have no idea of the date. It is not possible to tell that something is grue (on your own authority—you may of course be assured of it by an angel, etc.) without knowing the date. Perhaps it is logically possible that someone has sensory mechanisms that keep track of how much time has elapsed since the birth of Christ (or any other dating event). And these mechanisms may be crossed up with his colour perceptions in such a way that after t, newly discovered emeralds look different. He would also have to be sensitive to whether a thing has been previously observed, so that previously observed emeralds would continue looking like they used to look.

Such a person might well have severe problems getting his language in step with ours, if he had been one of us all along. It would take some ingenuity to work out the hypothesis that he was a person for whom 'grue' (or some translation of that term) was a purely qualitative predicate. Even if this is possible, its being logically possible for 'grue' to be purely qualitative does not make it purely qualitative. Furthermore, the bare logical possibility for colour predicates does not generalize to other predicates.

Consider the predicate 'dfalls',where X dfalls iff X is observed before t and falls or is not observed before t and doesn't fall. Then all suicide jumpers (without parachutes, etc.) have been observed to fall, but also, up to t, have been observed to dfall as well. Now it is after t and you are an addicted bungee jumper who is depressed at not being able to afford the bungee. Friends argue that you should just jump anyway, on the grounds that, like as not, you will just dfall, like everyone has done so far since time immemorial. Or again, the predicate 'dfirsts' meaning 'observed before t and does not finish first or not observed before t and does finish first'. All runners broken up in car crashes have failed to finish first in races held the next day. But they have also all dfirsted. So some follower of PI concludes that there is equal reason to believe that Boggs will dfirst as to believe that he will not finish first.

One difference between these B*'s and 'grue' is that they involve extensive consequences in a way that colour properties do not. It is not possible to make coherent a picture of someone for whom such predicates could qualify just as well as 'falls' or 'finishes first,' as not dependent on an arbitrary cut-off date. (No doubt 'winning' is not 'purely qualitative,' or 'observational', but the idea behind the objection about being purely qualitative was surely to complain that the B* predicates involved an arbitrary cut-off date.)

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

Goodman's position is that the absurdity of these projections is due to their involving predicates we have not developed the habit of applying in successful projections. It is of course true that we have not developed such habits. But it is perfectly possible logically that we should develop such habits without at all incurring the illusion that projecting these predicates after t would be anything but irrational. We might find that speaking in terms of 'dfalls' and 'dfirsts' keeps us on our toes about the date. The excitement might lead us to use such predicates heavily, in singular predictions about the behaviour of things we know will be observed safely before t. If we wanted universal laws independent of the time, the use of B* predicates would require skill in formulation. The use would be highly derivative and dependent on less esoteric terminology, but this is still compatible with getting habitual and successful use. Such habits would not give us any reason to expect that any patterns of projection would become reasonable which would not be reasonable without those habits (excepting of course predictions about what habits we will be exhibiting).

Russell held that the status of PI is problematic. Edwards and Will and Strawson present arguments with the upshot that the principle is true, even analytically true (though some seem not to recognize that this is what they are doing). They do this in spite of giving clear indications of rejecting the idea that induction needs or can be given a generic justification. That they none the less veer back to the generic is an indication that the tendency is deep rooted and difficult to curb. Generalized warnings against overgeneralization are apt not to be sufficient. Goodman offers an excellent showing that if the PI were a general truth, then in a perfectly possible case we could derive inconsistent results about what is justified. But then he concludes that the consequences of PI not being a general truth are about the ones Hume claimed. Here again is a mistaken estimate of what turns on the truth of PI.

If I were to toss a coin repeatedly and kept getting heads, I would no doubt be deeply impressed, and come to believe that there was something about the tossing setup that guaranteed heads. I would be mystified, and yet the pattern of habit formation would go to work, so that the coin's behaviour might come to seem another inexplicable commonplace. If asked to justify my expectations, I could offer nothing but what would sound like a case of PI. Whether or not I would be justified, I would certainly be excusable for forming the expectation I did form. This suggests that in some cases, mere past positives, etc. in the style of PI would have some weight in the bare way they are held to have weight in the PI. But to generalize from such a consideration is to risk losing sight of the rich variety

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press

of inductive reasoning. One may begin to see constant conjunctions cropping up through the variety of inductions and then vastly overestimate their importance in the critical assessment of the inductive performance. The tendency to think that one has discovered a test for the truth of claims to knowledge is endemic in Epistemology and provides occasions for critics of the subject to accuse its practitioners of delusions of grandeur. That branch of Epistemology which 'confines' itself to induction has been afflicted with some of the worst outbreaks of the tendency.

University of Virginia

https://doi.org/10.1017/S0031819198000205 Published online by Cambridge University Press