Episteme, 12, 2 (2015) 275–290 © Cambridge University Press
doi:10.1017/epi.2015.9

THE PROBLEM OF THE PROBLEM OF INDUCTION

ROGER WHITE
rog@MIT.edu

ABSTRACT

To solve the problem of induction we had first better know what it is. Some ways of formulating the worry about induction are underwhelming as they depend on assumptions that don't survive much scrutiny. Perhaps the most disturbing argument for inductive skepticism appeals to the claim that we could not possibly be justified in taking our inductive methods to be reliable independently of our use of those methods. And the use of inductive methods cannot give us justification to suppose that they are reliable. I argue for a new way to escape the first horn of this dilemma.

I'll try to say something new about the problem of induction. The problem is first to figure out what the problem of induction is. That's the problem of the problem of induction. Philosophers have probably had different things in mind by the problem of induction. For some the task is to try and show that reasoning inductively will lead us from true premises to true conclusions (for the most part), or perhaps showing that it is rational to proceed in this way. Depending on what counts as "showing" this might be a challenging job. And it might be that we are not up to the task. Perhaps there is nothing helpful to be said to someone who doesn't already find it obvious that we should reason inductively. That might be disappointing, but it is not clear that we should lose much sleep over it. Sometimes it is reasonable to proceed in a certain way without an argument as to why we should. What would be more challenging is if there is an argument for inductive skepticism – for the conclusion that we are not justified in drawing inductive conclusions – which proceeds from premises that strike us as highly plausible. The challenge will then be to explain where the argument goes wrong. I'll present what I take to be the strongest argument of this sort but argue that it has a significant weakness that has gone unnoticed.

So just what is the problem of induction? Presentations of the problem usually begin by noting that the premises of an inductive argument do not guarantee the truth of its conclusion. The sun may have risen every day in our past, but it is *possible* that this is so and yet the sun does not rise tomorrow. Introducing the problem in this can prompt a dismissive response: "That it is possible for the sun not to rise tomorrow only suggests that I should not be absolutely *certain* that it will rise. I can still be pretty damn sure, which is all we should ask for from an inductive argument" (see Stove 1973). The proponent of the problem usually responds by insisting that this objector has seriously underestimated the worry. What is at stake is not whether we can achieve rational certainty, but whether we should put any confidence in our inductive conclusions at all (Skyrms 1999). Why *should* we expect the sun to rise tomorrow, rather than expect the opposite, or remain noncommittal? But this kind of rhetorical question also invites a

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

straightforward response: "We should expect the sun to rise tomorrow because it has always risen in the past! Apparently you don't take this to be a very good reason. But why not? The fact that the conclusion doesn't strictly follow from the premise doesn't suggest that it is a poor reason unless we suppose that the only good reasons one can have to believe P are propositions that entail P. But why should we think that?"

The proponent of the problem of induction may now try a more subtle tack. "In drawing and in defending your conclusion that the sun will rise tomorrow you implicitly followed a certain rule inference, very roughly: expect things to continue in the same salient patterns that we have observed. If you are justified in your conclusion you ought to be able to justify your use of this kind of rule. But this we cannot do without objectionable circularity (to argue that the rule has served us well thus far is just to appeal to the same rule, which is inappropriate when the status of that very rule is what is at stake). Heroic attempts to justify our inductive practices have failed miserably. So we are in an epistemological pickle."

While this gets closer to the heart of the problem, it is far from rigorous. There are at least two distinctions that we should make when considering this kind of argument. The first is between the *status* of justification and the *activity* of justification (Alston 1989a: 81–114). To judge that someone's belief that P is *justified* is to positively evaluate the belief and ascribe it a certain epistemic status. To *justify* a belief is to engage in the activity of defending it by offering reasons in support of it.

What is ultimately at stake in the problem of induction is the *status* of our inductively supported convictions, such as whether we are justified in expecting the sun to rise tomorrow. Yet in many presentations of the problem of induction, the skeptical worry is supposed to somehow arise out of our failure at the task of justifying something (our conclusion, our use of certain rules or inferential methods, an allegedly crucial underlying assumption like the Uniformity of Nature). But it is far from obvious what the relation is between the activity of justification and the status of being justified. In addressing the question of whether I am justified in believing P, it might seem worthwhile to consider what I might say in defense of P. For if I can articulate some good reasons I have for believing P, this could satisfy the worries I had about the epistemic status of my belief. But is it *necessary* to be able to articulate such reasons? In many cases we feel satisfied that we have adequate reasons for our views, although we are unable to say what they are (e.g., when we "read" someone's facial expression, judging that she is upset. We can't begin to say what the subtle visual cues are that lead us to think this). Other things just strike us as obvious, such as that bananas are bananas – no fancy defense is required, nor would one add much to the status of our conviction.

If there are some more theoretical grounds for supposing that being justified in one's beliefs somehow requires giving or being able to provide a justification for them this remains to be seen. But there are reasons to be skeptical. The temptation to suppose that being justified in a belief requires being able to give a justification for it seems to derive from certain dialectical situations. If someone accuses me of being unjustified in holding a certain belief, I might feel obliged to respond by giving a defense of it. If I succeed my critic will come to share my belief and so most likely come to agree that I was justified after all. But my success in this endeavor, and even the appropriate way to proceed depends on certain factors relative to my audience. For instance, it is considered inappropriate for me to appeal to premises that I know my opponent will reject. In the context of skeptical challenges, we are often told that we must avoid "begging the

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

question against the skeptic". It is quite unclear what this amounts to. But whatever it is, it seems rather odd that the status of my beliefs should depend in any way on my ability to appropriately and successfully engage in a debate with any particular kind of subject. Surely the epistemic status of my beliefs hinges on factors about *me* and my beliefs, experiences and the like, not on who I happen to be talking to.

This is especially so in the case of fundamental convictions like the status of induction. It really just strikes me as obvious that the right conclusion to draw from the fact that the sun has risen every day is that it will rise again tomorrow. My conviction is not based on any fancy argument, let alone one aimed at persuading a skeptic. I think there would be something wrong with anyone who didn't share my conviction. If, as I think, inductive skeptics are irrational, then it's no surprise that I can't come up with an argument that could bring them around. I can't persuade the tinfoil-hat guy in the park that aliens are not trying to tap into his brain. But that doesn't reveal much about the status of my own beliefs. Of course it is possible that my complacent trust in induction is misguided, and perhaps I ought to be open to this. But no serious doubt has been raised by my failure at heroic tasks of justification.

The second reason to be suspicious of sketchy presentations of the problem of induction is that there is a danger of falling into epistemological *level-confusions*. It is one thing to be justified in believing P, another to be justified in believing that I am justified in believing P, and so on. Many philosophers appear to have been guilty of conflating or sliding between first and second or higher-order questions of justification (see Alston 1989b: 153–71). As we've noted, if asked to justify my expectation that the sun will rise tomorrow, I might very well simply reply that it has risen every day in the past. In motivating the problem of induction, the skeptic has asked me to defend my use of inductive rules to arrive at this conclusion. What the skeptic is questioning is not directly whether the sun will rise tomorrow, but whether I am justified in concluding that it will. The apparently difficult task of justification is that of arguing that I am indeed justified in this conviction. But then even if we could legitimately move from my failure to justify a belief in P to my not being justified in believing P, this would immediately get us only as far as the conclusion that I am not justified in supposing that I am justified in believing that the sun will rise tomorrow. But it is the first-order justification that the problem of induction is ultimately concerned with. Of course there may well be principled ways to move between epistemic levels in order to arrive at a disturbing first-order conclusion. But these are subtle and controversial matters that must be handled with care (Horowitz 2014).

I've suggested that we ought to be dubious of the usual hand-waving presentations of the problem of induction. It's time to explore how to put together a stronger and more rigorous case for inductive skepticism. I've noted that as Alston (1989b) has argued, philosophers seem prone to conflate or slide between epistemic levels. Perhaps a more charitable view is that we have an implicit grasp of certain principles linking different levels. In any event, it is worth considering what such principles might be and if they are correct. The sort of principle that we need to take us from higher-level to lower-level skepticism is an Epistemic Ascent principle.

Epistemic Ascent: If S is justified in believing P, then she is justified in believing that she is justified in believing P.

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

The principle says roughly that matters of justification are epistemically available to me. If I'm justified then I ought to be able to tell that I am. Given such a principle, from the claim that I am not justified in believing that I am justified in believing that the sun will rise tomorrow – because I'm not justified in believing that I have followed the appropriate rule in drawing this conclusion – we can conclude that I am not justified in expecting the sun to rise.

The principle will need some refinements to be a serious candidate. As it stands it suggests that having justification entails justification at an infinite hierarchy of levels: justification that one is justified in believing that one is justified in believing that one is justified . . . . A more modest principle that might serve the skeptic's purposes is something like this:

> **Epistemic Ascent***: If S considers whether she is justified in believing P, and she is not justified in believing that she is justified in believing P, then S is not justified in believing P.

This formulation sidesteps the question of how we are to evaluate the fellow who doesn't even consider epistemic questions when forming his beliefs. The plausibility of the principle derives from the regulative role that justification should play. We seek justification as a means to pursuing the truth. In an earnest inquiry into whether P, I might ask myself whether I would be justified in believing it given what I have to go on. My answer should guide me in my beliefs. If considerations of justification did not play this regulative role it might be hard to see why we should care about it.

So suppose that I consider whether I am justified in believing P, and as it happens I am not justified in believing that I am so justified. The following seems like an appropriate line of thought in this situation. "Maybe I'm justified but maybe I'm not. If I am I'm likely to be right. If I'm not justified then I could be right just by dumb luck, but this is uncommon. So since I'm unsure about whether I'm justified in believing P, I'm unsure whether my belief in P is correct." In other words, if I'm not justified in believing that I'm justified in believing P then reflection on the matter should lead be to doubt P, i.e., I'm not justified in believing it.

As plausible as this principle may be, it will only take us to inductive skepticism if we can establish second-order skepticism: skepticism concerning our justification for drawing inductive conclusions. But is not easy to see how to do this. I've already complained that demands to justify a belief don't cut much ice. For many, the skeptical worry arises from the apparent contingency of our success at inductive inference. In order for us to reliably draw true conclusions by the application of inductive rules, the world has to cooperate. If the world happens to be a nice orderly place, then our inductive methods will serve us well, but who's to say that it is? Sure, those parts that we have actually observed have fortunately been quite regular. But it is the wider-scale Uniformity of Nature on which the success of our inductive practices depends. And we have no way of telling whether the world is conducive to the successful application of inductive methods.

But while these kinds of points may seem disturbing, they can also seem irrelevant to questions of justification. Whether or not I am justified in drawing an inductive conclusion seems to depend solely on what the conclusion is, and how I drew the conclusions from what assumptions. Whether or not I'm justified in inferring that the sun will rise tomorrow from its having risen in the past shouldn't depend on such contingent matters as whether the future tends to resemble the past. What matters is whether my premises support my conclusion, and this seems to be a non-contingent matter. But if the epistemic status of my belief

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

does not hinge on such contingent, external matters that are beyond my ken, what is to stop me figuring out whether my beliefs are justified from the armchair? Careful a priori reflection reveals that, other things being equal, you are justified in believing that something is red if it appears red to you (even if unbeknown to you, you happen to be a brain in a vat). Similarly, careful thought on the matter reveals roughly that other things being equal, you are justified in expecting certain salient patterns to continue (even if unbeknown to you, the world is an irregular place, which will trip up the inductive reasoner at every turn). Of course one might also be skeptical that a priori reasoning can reveal to us even the truth of non-contingent epistemic matters. But the case for inductive skepticism will not be very impressive if it hinges on another controversial form of skepticism.

The suggestion above might be cold comfort to those who are impressed by the problem of induction. They are apt to complain that it is all very well to smugly insist that we can tell a priori that we are *justified* in our inductive conclusions (regardless of whether they are mostly, or even ever correct!). But what we really care about is the *truth*. What good does it do me to know that my inductive conclusions are *justified* if for all I can tell induction might be leading me hopelessly into error? (Compare: What good does it do me to know that a brain in a vat is justified in his perceptually based beliefs, when for all I can tell this is *my* predicament?)

The non-skeptic might respond as follows. I care about truth as much as you do. But there are really no grounds for doubting that we have obtained it by our inductive methods. Quite the opposite. A priori reflection reveals that we are indeed justified in believing that the sun will rise tomorrow, that bread will continue to nourish us, and a whole host of other inductive conclusions. If these inductive conclusions are true, it follows that induction is not leading us hopelessly into error. While it might fail us now and then, it is typically leading us to the truth. So if we are justified in believing these inductive conclusions then we are justified in believing that induction is not typically leading us into error.

I think the skeptic can do better at pressing his concern. He could appeal to the following principle.

> **Reliability Principle**: If S considers whether the methods or rules she followed in concluding P are reliable, and she is not justified in believing that they are reliable, then she is not justified in believing P.

The Reliability Principle draws a link between justification and reliability, but not the controversial one that justification *requires* reliability. Instead it imposes a second-order constraint on justification. Like the previous principle, this one leaves open what we are to say about those epistemological innocents who form their beliefs without considering the reliability of their methods. So the principle will not lead to general inductive skepticism regarding the man on the street. But you and I are not so innocent (or at least you won't be when I'm through with you). We are forced to consider the reliability of our inferential rules. Combined with the assumption that we are not justified in accepting the reliability of our inductive rules, we get a skeptical conclusion concerning ourselves, and that is more than bad enough.

Why accept the Reliability Principle? It depends on the plausible idea that our judgments about the reliability of our methods should restrain our use of these methods. Suppose I consider whether my inferential methods are reliable, but I am not justified in believing that they are. Then clearly I should not believe that they are. I shouldn't

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

disbelieve their reliability either. I should simply remain undecided as to whether my methods are reliable. But if I don't think my methods are reliable, wouldn't it be foolish to accept the results of those methods?

Although this seems very plausible, I think many might end up biting the bullet and denying it, if only as a way of avoiding a skeptical conclusion. So it is worth driving home just how difficult it is to deny this. Start with some analogies. I have a pocket calculator that I've used to conclude that 20826/89 = 234. I now consider whether the calculator reliably gives the right answers. For whatever reason, I'm in no position to tell whether it does or doesn't. It would be crazy to continue to believe that 20826/89 = 234. I don't think anyone would want to deny this. The same applies to my use of, say, a long division algorithm. I couldn't sensibly have any confidence in my answer if I couldn't tell whether the rule that I used at least typically gives the right answer. Shouldn't this apply equally to the use of inductive rules?

As plausible as this seems, many philosophers do explicitly make an exception in the case of perception. According to the position which has become known as Dogmatism, for certain contents P, I am justified in believing P on the basis of its appearing to me that P, even if I am not justified in believing that my perceptual faculties are reliable (provided I have no reason to judge them *un*reliable) (Pryor 2000). According to the Dogmatist, the rule sanctioning transitions from experience to belief has a special status. It just is appropriate to follow this rule regardless of what justification we might have for taking it to be reliable (at least provided we have no reason to suppose that it is unreliable). While the case of inductive rules is not so commonly discussed, it might be natural to extend the dogmatist thesis to cover induction and other inferential rules. We might suggest that the very most *basic* epistemic rules enjoy this special status. Rules that are derived from others like the long division rule must earn their keep by one's being justified in believing that they are reliable (at least insofar as you consider the matter).

I think this kind of view has many attractions, not least of which is that it allows us an escape route from inductive skepticism. However it faces serious difficulties, as I will argue now. Let's suppose that prior to forming an opinion as to whether P, I consider what kind of tendency the method I am planning to use to answer this question has to yield true beliefs. But unfortunately in this case I am in no position to tell that it is reasonably reliable at yielding the truth (even if in fact it is reliable). Now reliability here is a modal notion that does not strictly imply anything about actual success. An unreliable method might get lucky and work every time; a reliable but fallible method might consistently fail. Reliability is just a matter of a *tendency* to produce a certain result. But there are obvious epistemic links between judgments about tendencies and expectations of outcomes. Suppose I can't tell what kind of coin I'm holding. Perhaps it has a strong tendency to land heads. Or perhaps it is biased toward tails. Or maybe it's an ordinary fair coin. I can't tell. It would be foolish in this situation to have much confidence that the coin will land heads (unless perhaps if I was privy to some special information, like if I could "see the future"). This obvious fact can be supported by what Lewis called the Principal Principle (1980): (roughly) if P is a possible outcome of a future chance event, one's subjective probability for P, conditional on the objective chance of P being x, should be x. So my confidence in P should be a weighted sum of the possible objective probabilities of P (each weighted by my confidence that that is P's objective probability).

So given that I have no justified opinion as to the reliability of the method I plan to employ to determine whether P, I should likewise have no idea as to whether I will end

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

up with a true belief on the matter. It would be very odd to think, "I haven't any idea whether this method tends to yield the truth, but I'm quite confident that it will in fact do so nevertheless." If the method does happen to be reliable, then sure, it is to be expected that it will succeed. But for all I can tell it won't, in which case it will be sheer luck if I get the right answer. Furthermore, just suppose that I could somehow reasonably expect actual success from the method whose reliability I can't now tell (perhaps a prophet who obtains information from the future informs me that whether by reliable means or by lucky fluke, my method will in fact yield the correct answers). If I plan to use the method repeatedly, then I will be justified in expecting it to have a long run of successes. Quite a remarkable feat for a method if it is unreliable! If a prophet were to reveal to us that a certain coin will in fact land heads every time, the only sensible thing to conclude would be that the coin is heavily biased and so had a strong tendency to land heads. All in all, it seems that we cannot rationally combine agnosticism about the reliability of a method, with confidence that it will nevertheless be successful.

At this point it seems that I shouldn't proceed with the application of this method. For my goal is to obtain the truth. And I shouldn't proceed in a way that I don't expect to be conducive to achieving this goal. Suppose I go ahead and apply the method anyway, forming the belief that P. Surely reflection on the fact that I formed the belief in a way that was inappropriate (and I have gained no better reason since) should lead me to lose all confidence in my conclusion. I cannot coherently combine my doubts that the method will yield the truth with my newfound confidence in the conclusion P. For I now know that the method has resulted in my believing P, and in conjunction with P this entails that the method yielded truth. Something has to give, and the appropriate response is surely to lose my confidence in P. In other words, one can't rationally combine doubts about the reliability of the method by which one came to believe P, with confidence in P. Since one ought to doubt the reliability of one's methods if one is not justified in believing them to be reliable, we have arrived at the Reliability Principle.

Someone might try to resist this last step by insisting that for some methods (and this includes our most fundamental inductive methods) even if you shouldn't have expected the method to yield the truth due to your inability to tell that it was reliable, once you have applied the method and formed the belief, you are entirely justified in maintaining that conviction. For it is in fact the *right* method to follow, and this is ultimately what matters for justification. Call this position induction-dogmatism.

Here are a few ways of drawing out odd consequences from this kind of position, which are modified from some previous work (White 2006). For each argument I will be assuming that prior to applying the method I am not justified in believing that it will yield a correct answer (because I am not justified in believing that it is reliable). First, suppose that I am justified in accepting induction-dogmatism. Then before applying the method, I am justified in believing that I will soon be justified in believing that the method was successful, without my having any cognitive mishap or gaining any evidence in the meantime that I can now judge to be misleading. For I am justified in believing that I will soon have a justified belief in the answer given by the method. And from this together with my knowledge that the method has given this answer I can infer that the method was successful. But other things being equal, if I am justified in expecting to have a justified belief that P, then I am already now justified in believing that P. Why must I wait around for this justified belief to arrive before I can reasonably conclude that P? So it follows that prior to using the method I *am* justified in expecting it to yield a true belief. But this

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

contradicts our earlier conclusion that I am not so justified when I am not justified in believing the method to be reliable. So I can't be justified in believing induction-dogmatism after all.

Second, suppose I were to use the method repeatedly to form numerous beliefs with contents P1, P2, . . . , I will be able to tell that I have formed these beliefs. So from a proposition of the form

> P1 and I believe P1, P2 and I believe P2, . . . , and these are all the beliefs that I have formed by this method

I will be able to validly deduce that my method has successfully yielded true beliefs every time that I have used it. But as I suggested above, the only reasonable conclusion to draw would be that my method was reliable. For it would be a bizarre coincidence if an unreliable method was successful every time. What else could plausibly explain the method's actual success if not its general reliability? So induction-dogmatism endorses "bootstrapping" (Fumerton 1985, Vogel 2000). Beginning with doubts about the reliability of my method, I can simply use the method itself to confirm that it is indeed reliable! But this is absurd. Surely the only way that one can reasonably confirm the reliability of an inferential method is to use some independent means of assessing its results.

For the final objection to induction-dogmatism consider the following propositions concerning the use of a method to figure out whether P.

> **MP**: The method gives the answer P.
> **MFP**: The method falsely gives the answer P.

If prior to employing the method I am not justified in believing that the method will yield the correct answer, then at least one of either proposition P or not-P is such that I am not justified in particular in expecting the method not to falsely give that as an answer. For if the method will not falsely give either answer, then it will give the correct answer if any. So if I was justified in believing both that it will not falsely give the answer P and that it will not falsely give the answer not-P, then I could tell that it will give the correct answer. So let's suppose that I am not justified in particular in expecting it not to falsely give the answer P, i.e. I'm not justified in believing not-MFP. But according to induction-dogmatism, if I do come to believe P by use of the method then my belief in P will be justified. But P is inconsistent with MFP. So if I can be reasonably confident that P, then I can be reasonably confident that not-MFP. So as a result of using the method, I will have rationally gained confidence in not-MFP.

Already this seems odd. At first I think to myself, "Will this method falsely give the answer P? I'm not at all sure." I now use the method and note that it gives the answer P. "Ah, good. Luckily it didn't falsely give the answer P." But we can strengthen this point. When I use the method and note that it gives the answer P, I learn MP. Now MFP entails MP. So P(MP|MFP) > P(MP) and hence P(MFP|MP) > P(MFP), and hence P(not-MFP|MP) < P(not-MFP). The method's giving the answer P *lowers* the probability that it doesn't falsely give the answer P. So if our degrees of confidence are to conform to these probabilities, then my confidence in not-MFP should *drop* whereas induction-dogmatism entails that it should *rise*.

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

I've been arguing that the Reliability Principle is very hard to deny. Back now to inductive skepticism. What we need to reach a skeptical conclusion is the premise that we are not justified in believing that our inductive methods have a tendency to get to the truth. But how might we argue for this? Not by demanding that we justify our belief that induction is reliable in order to be justified in believing it. A more tempting suggestion was touched on above. It is commonly thought that the reliability of our inductive methods is a radically contingent matter. If we are fortunate enough to inhabit a very orderly world – one with pervasive and widespread regularities of the kind that our inductive rules sanction projecting – then induction will serve us well, for the most part taking us from true premises to true conclusions. But the world needn't be this way at all. It could be a big mess. There are possible worlds in which people earnestly follow inductive rules as we do but are led hopelessly astray. Patterns which emerge in their data turn out not to extend to unobserved cases. They continually find just enough Fs that are G to inductively infer that the next F will be G only to find that it is not. These worlds are induction-unfriendly, in the way that Matrix worlds are perception-unfriendly. How then are we to tell that ours is not one of them?

The worry here about induction can begin to seem even more pressing than those about perception. For it seems that we needn't invoke anything as outlandish as powerful deceptive demons, or scientists wiring up brains in vats. It is easy to feel dismissive of such bizarre possibilities. For induction to be reliable it seems that we need the world to be structured in a certain orderly way. But would it be all that strange for it not to be this way?

To raise doubts about our justification for believing that our inductive methods are reliable, we needn't depend on challenges to defend the claim that they are. General considerations about what it takes to be justified suggest that the reliability of induction is just not the kind of thing that one could be in a position to have an opinion about. The reliability of induction apparently depends on contingent facts about the large-scale structure of the world. But it can seem pretty far-fetched to suppose that we can tell what the world is like in the relevant respects from the armchair. Imagine someone claiming that they could tell that there are currently an odd number of bananas in Costa Rica. "It's just obvious" he insists. This sort of thing is just not obvious and can't be. Similarly, when thinking of the reliability of inductive rules in terms of what the world would have to be like to be conducive to inductive success, it can seem far from obvious that it is this way. But even if we fancied that we had a deep philosophical insight that the world was induction-friendly, or even a subtle a priori proof, on reflection we ought to be suspicious of this (as we would of any alleged philosophical grounds for supposing that there are an odd number of bananas in Costa Rica). The general principle seems to be that we are in no position to tell a priori, the truth concerning contingent matters external to ourselves.

Is there any other way that we could tell that our world is conducive to the successful application of inductive rules? We haven't been able to survey the world past, present, and future and discover that nature is fairly uniform. Those parts of the world that have actually been observed have been nice and orderly. This might explain the remarkable success that our inductive methods have enjoyed thus far. But the general reliability of induction requires regularity on the broader scale including the future, the past, and other regions that we haven't seen yet but might. It appears that the only way we could tell how orderly the world is on a broader scale is by inductive inference from what we have observed.

The standard objection to this suggestion is of course that it is illegitimate because it is *circular* – we are using a method to defend the reliability of that very method. But this

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

complaint is not very helpful as it stands. Just what circularity is and what is wrong with it turns out to be a surprisingly subtle matter. The topic is slippery enough that several philosophers have given defenses of a kind of 'inductive defense of induction' that cannot be easily dismissed (Black 1974; van Cleve 1984; Papineau 1993). It may be obviously dialectically inappropriate to try to *persuade* someone that the world is induction-friendly by using induction. But the person who suggests that we know by induction that the world is orderly needn't be understood as doing any such thing. He doesn't think that the reliability of induction should be in question in the first place. He is just suggesting how it is that we can tell (as of course we can?) that the world is induction-friendly. Besides, there are plenty of legitimate lines of reasoning that fall broadly under the description of "using a method to defend the reliability of that very method." For example I can listen to my doctor's positive report of my hearing test, and thus boost my confidence that my auditory faculties are reliable. Accusations of circularity must be handled with care.

Nevertheless, given other assumptions in our argument thus far, I think we can see why the suggestion that we know inductively that the world is induction-friendly will not work. I have argued that in order to be justified in drawing an inductive conclusion, one must be justified in believing that induction is reliable (at least if one has considered the matter). We are now considering whether induction is reliable. So in order to be justified in inductively concluding that the world is induction-friendly from the premise that the observed part of the world is induction-friendly, we must already be justified in believing that the world is induction-friendly. So we are in a Catch-22: we can't meet the necessary condition on being justified in inductively concluding that the world is orderly, since already being justified in believing that the world is orderly is required in order to meet the necessary condition for inductively concluding that the world is orderly.

So it appears that we simply can't be justified in believing that induction is reliable. And together with the Reliability Principle, this entails that those of us who have considered whether our inductive methods are reliable are not justified in expecting the sun to rise tomorrow. With the parts in place, we can now set out the structure of the argument.

P1. Reliability Principle: If S considers whether the methods or rules she followed in concluding P are reliable, and she is not justified in believing that they are reliable, then she is not justified in believing P.

P2. The reliability of our inductive methods is a contingent matter of the large-scale orderliness of the world.

P3. One cannot have a priori justification for believing a contingent proposition concerning matters external to one's self.

P4. The only way that experience could justify us in believing that the world is orderly, is if we were justified in inductively inferring its orderliness from its observed orderliness.

P5. There is no other way for us to be justified in believing that our inductive methods are reliable.

C. Therefore, those of us who have considered whether our inductive methods are reliable are not justified in holding any inductively formed belief.

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

The argument no doubt needs some refining, but I think it makes about as strong a case as we can find for inductive skepticism. Of course now that we have set it out there are going to be plenty of possible lines of resistance to consider. I am going to focus on one that I think has been overlooked. I have been claiming that it is easy to see how our inductive methods could have been quite unreliable, and hence that we are lucky at best if we inhabit an inductive-friendly world. This is what I want to start to question. (This is not actually one of the argument's premises, so questioning it won't directly show us where the argument goes wrong. But we'll get back to that).

First let's note that inductive conclusions can be more or less ambitious. From data consisting of Fs that are G we might conclude that the next F will be G, or that most Fs in the vicinity are G, or that all Fs are G. Obviously we should put less confidence in the stronger theses, and it is less disturbing that we should often be wrong in our more extravagant predictions. Our data do not support the conclusion that the sun will *always* rise each day – we know that it won't. And our confidence that it will continue to rise for some time depends more on knowledge of physics than a simple extrapolation from observed sun-risings. To avoid complicated issues, I will restrict my focus to the most modest inductive inferences, those that take us just to the next instance of some regularity. After all, if the argument for inductive skepticism is successful, it should entail that *any* extrapolation from data is unjustified.

Let's think about what it would take for a world to be induction-unfriendly. We are apprentice demons given the task of constructing a world in which inductive reasoners will tend to form false beliefs. As a beginner's task we are given just one binary parameter to set in a sequence, such as the sun's rising or not rising on consecutive days. In our plan we will represent risings as '1's and non-risings as '0's. Now obviously the first thing that we want to avoid is a sequence displaying a simple exceptionless regularity like

11111111111111111111111111111111111111111111111111111111111111

or even something fancier like

00111001110011100111001110011100111001110011100111001110011100

These sequences like many others are perfect for induction. Inductive reasoners will latch onto these patterns, predict that they will continue, and be right. So perhaps then we should try a disordered sequence like

01100011110010011111010101001001111000011010101110010001111010101010

But while this may not be the best sequence on which to apply inductive rules, it is not one that will tend to lead inductive reasoners into *error*. For induction is barely *applicable* to this sequence. In order for our standard rules of induction to be applied, our data must contain a salient regularity of a reasonable length. The only opportunity that we could have to apply inductive rules to this sequence would be to a second-order regularity: All of the larger subsequences that we observe are disordered within. Hence we could inductively infer that the sun will continue to rise and not rise in a disorderly fashion. And we would be correct. So really this is a world in which inductive rules are very reliable at leading from true premises to true conclusions.

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

What we will need to do to lead our victims into error is give them enough of a pattern to license an extrapolation to further instances, and then switch things to make their conclusion incorrect. Let's try

IIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIOOOOOOOOOOOOOOOOOOOOOOOOOOOOOO

On the thirty-second day of this sequence our victims might inductively infer that the sun will rise the next day, and they will be mistaken. That's a small victory for us. But unfortunately the vast majority of opportunities that they will have to apply inductive rules are ones in which they will get it right. Most of the uniform subsequences here are followed by another day of the same. So the truth to error ratio among inductive inferences will tend to be quite high.

We will have to do better than this. The best we can do, it seems, is to make the "switcheroos" more frequent in order to catch out our inductive victims as much as possible. How about

IIIIIOOOOOIIIIIOOOOOIIIIIOOOOOIIIIIOOOOOIIIIIOOOOOIIIIIOOOOO

On every 5th day the inductive reasoners may form a false expectation of what the sun will do next. But the trouble is that our inductive rules only sanction very tentative conclusions from data consisting of only five instances of a pattern. So it is not very alarming that such feeble inferences should often be false. We face a dilemma here. The shorter we make the uniform subsequences, the more tentative the inferences that are sanctioned by our inductive rules. But the longer we make these subsequences, the more opportunities we provide our victims to make successful inductive inferences. It is surprisingly difficult to maximize the confidence sanctioned by inductive rules applied to the world while minimizing the reliability of those rules in the world. Furthermore, it is only those inductive reasoners who happen to observe just a uniform segment of the sequence that will be led into error. Anyone who observes some longer sequence will see the sun rise some days but not others and so will not be in a position to tell what it will do next. They might observe a long enough sequence to note that five risings alternate with five non-risings, in this case they will be able to make confident predictions of what the sun will do next. And their predictions will be accurate.

Our task is turning out to be much more difficult that it first seemed. It is worth noting that this is not so for other possible inferential rules. Consider counter-induction, which tells us to expect the opposite of what has come before. The counter-inductivist predicts that the sun will not rise tomorrow, given that it has always risen in the past. The possibility of counter-induction is often invoked in presentations of the problem of induction with the challenge to say what could possibly make one the right rule to follow rather than the other. But unlike induction, it is not hard at all to cook up a world in which counter-induction is hopelessly unreliable. We just have to make the world very regular. If the sun always rises then the counter-inductive reasoner will continually be expecting it not to rise the next day, with ever-increasing confidence. And the repeated failure of his counter-inductive rule to yield the truth, he will counter-inductively increase his conviction that it will succeed the next time. It is actually quite hard to imagine a world in which counter-induction *is* reliable. Not only is it difficult to find a general structure conducive to counter-inductive success (corresponding to the difficulty of constructing a world

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

conducive to inductive failure), any initial run of counter-inductive success would appear to be short-lived. For a second-order application of the counter-inductive rule would then entail that first-order applications will now fail, so one should not make the opposite prediction.

It might be tempting to conclude that there is no such thing as an induction-unfriendly sequence. That would be a nifty result indeed. It would mean that *necessarily* induction is reliable on a binary sequence, and hence that P2 is false, at least with respect to such sequences. But I doubt that this can be right. Part of the difficulty here is that it is hard to say in detail just what our inductive practices are exactly. But let's just say that our inductive methods can be captured in a function I(S) mapping initial sequences to verdicts of the form *infer 0 next*, *infer 1 next*, or, *withhold judgment*. We ought to be able to use this function to construct a sequence such that most of the time it tells us not to draw any inference. But on the occasions in which it does tell you to infer the next digit it makes the wrong prediction. So it seems it has to be at least possible for induction to be unreliable on a sequence. The point however, is that this is a specific sequence constructed in a special way. We'll come back to this.

These difficulties for the demon only get worse when we consider not just a predetermined sequence of the sun's behavior, but instead the overall distribution of properties of objects. Our next task is to assign colors to the emeralds in a way that maximizes our victims' chances of making false inductive predictions. In order to give them a decent chance of getting data that license any inductive inference at all, we will have to make a fair number of the emeralds the same color, say green. Of course it will not do to make all or even most of them green, for then the application of inductive rules is very likely to be successful. For in this case we can expect them to observe a lot of green ones, predict that they will see another green one next, and most likely be right. Perhaps we should make half of them green and the rest blue. Now of course it is possible for our victims to select a lot of green ones, predict that the next one will be green but get a blue one instead. But it is not enough just to make this *possible*. We are looking to make a world in which the application of inductive rules is *unreliable*. We need our victims to *tend* to fall into error, not just to do so if they are very unlucky. But it is a mathematical fact that if half of the emeralds are green, then the vast majority of subsets of the emeralds large enough to sanction an inductive inference will also contain approximately half green and half blue. So unless our inductive reasoners have a tendency to examine some emeralds over others, then they will have a strong tendency to observe about as many green as blue emeralds. From these data they cannot inductively predict what color the next emerald will be. However they can infer that they will continue to find about as many blue as green emeralds and that about half of the total emerald population is green and half blue. And in this prediction they would be right.

It looks as though we are going to have to take a more aggressive approach. Rather than just assigning properties to things and leaving it to chance how our inductivist victims will obtain their data, we must somehow encourage them to obtain *misleading* data. Given sufficient demonic powers perhaps we can predict which objects they will observe when, and carefully arrange properties to maximize opportunities for unsuccessful prediction. Or perhaps we can surreptitiously guide our victims to observe unrepresentative samples of things. This effort will probably require a large team of demons working around the clock.

But we should note that there is a cost to the skeptical dialectic in invoking the more aggressive demon hypothesis. I noted earlier that the skeptical worry about induction

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

can seem particularly strong because only very modest suppositions are required to get a situation in which induction is unreliable. It had seemed as though the world had to be a very special way in order for our inductive methods to tend to yield the truth. Our talk of the demon making an inductivist-deceiving world was just a way of thinking about what an inductive-unfriendly world would have to look like, with or without a demon at work. But now it looks as though we have to consider very extravagant hypotheses to get cases in which induction is unreliable, ones that essentially involve something like a powerful intelligent agent at work with the specific purpose of leading us astray.

The importance of this comes out in connection with P3 of our skeptical argument:

> **P3**. One cannot have a priori justification for believing a contingent proposition concerning matters external to one's self.

This doesn't seem quite right. Consider the following propositions.

1. The number of electrons in the universe is not prime.
2. My descendents will not all have the same birthday.
3. The stars in nearby galaxies are not arranged to form a portrait of me.

All of these are contingent and concern matters external to me. Yet it seems entirely reasonable to have a high degree of confidence in each of them. It is not easy to say exactly why, but it has something to do with the controversial Principle of Indifference. It is notoriously difficult to state what the relevant rule is in a way that is informative enough but doesn't lead to contradictions (van Fraassen 1989; White 2010). Nevertheless, it is hard to deny that something along these lines is correct. In a state of ignorance concerning a class of exclusive possibilities, where we can see no reason to expect one to be true rather than another, we should put about equal credence in each. When we consider the vast array of ways that stars can be arranged, it would be unreasonably arbitrary to put a lot of confidence in some possible arrangements over others. Star arrangements that compose a portrait of me make up a miniscule proportion of the total possibilities. So we should put very little confidence in that obtaining, and hence be quite sure that it won't.

So if a situation in which induction was genuinely unreliable required a *very specific* arrangement of things, then perhaps we can be a priori justified in believing that this is not the case. Of course we can't entirely rule it out. But it should seem far less disturbing.

It would be an exaggeration to suggest that we have the makings here of an a priori proof that we are justified in believing that our inductive methods are reliable. But my concern here has not been to show that inductive skepticism is false (we knew that already), but rather to see how compelling a case can be made *for* skepticism. And here we see at least one weakness. P3 will need some serious reworking to be plausible. And it is not clear how that can be done. Here is a kind of weakening of P3 that makes it more plausible and avoids the counterexamples above.

> **P3*** Let {P1, P2, . . . , Pn} be a partition of the region of logical space not ruled out by our evidence and such that we have no more reason to suppose that any one element of the partition obtains over another. If P is true in a large proportion of these possibilities then we are not justified in having much confidence in P.

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

But it is not at all obvious that the proposition that induction is unreliable takes up a large proportion of any such partition.

Let me end with one last worry about where this has got us. Suppose the Epistemology Oracle is kind enough to inform us that induction – by which she means our ordinary cannons of inductive reasoning – is indeed reliable. Not perfectly reliable of course. But use of induction does tend to lead us from true premises to true conclusions for the most part. This sounds like great news. We can trust the Oracle. Our worries about induction are over. Now we *know* that induction is reliable how could we have any qualms about applying it?

Not so fast, says the skeptic. Consider *schminduction*. Schminduction is a rule that coincides with induction in all but one possible case, which happens to be the case we are currently in. That is, given the actual data we have, schminduction instructs us to conclude that the sun will *not* rise tomorrow. For any other possible data and conclusion, schminduction is just like induction. Why then apply *induction* rather than *schminduction* the skeptic asks? It won't do to answer, "because we know that induction is reliable." If this is so then we also know that schminduction is reliable. The question is whether induction is *more* reliable that schminduction. If I have no idea which of these is more reliable then I can hardly conclude that the sun will rise tomorrow. So the skeptic give us a new argument with the following variant of P1:

> P1*: In order to be justified in concluding that the sun will rise tomorrow we must be justified in believing that induction is more reliable than schminduction.

And he will go on to claim that we are in no position to be justified in this comparative reliability claim. But note that given our evidence, whether induction is more reliable than induction just is a matter of whether the sun will rise tomorrow. So the premise amounts to the claim that in order to be justified in concluding that the sun will rise tomorrow we must be justified in thinking that the sun will rise tomorrow. This is right but doesn't make for much of a skeptical argument. After all, we are justified in concluding that the sun will rise tomorrow because we know it has always risen in the past.[1]

REFERENCES

Alston, W. 1989a. 'Concepts of Epistemic Justification.' In *Epistemic Justification: Essays in the Theory of Knowledge*, pp. 81–114. Ithaca, NY: Cornell University Press.

—— 1989b. 'Level Confusions in Epistemology.' In *Epistemic Justification: Essays in the Theory of Knowledge*, pp. 153–71. Ithaca, NY: Cornell University Press.

Black, M. 1958. 'Self-supporting Inductive Arguments.' *Journal of Philosophy*, 55: 718–25.

Fumerton, R. 1985. *Metaepistemology and Skepticism*. Lanham, MD: Rowman and Littlefield.

Horowitz, S. 2014. 'Epistemic Akrasia.' *Noûs*, 48: 718–44.

Lewis, D. 1980. 'A Subjectivist's Guide to Objective Chance.' In R. Jeffrey (ed.), *Studies in Logic and Probability*, pp. 83–132. Berkeley, CA: University of California Press.

Papineau, D. 1993. 'Reliabilism, Induction, and Skepticism.' *Philosophical Quarterly*, 42: 1–20.

1 Thanks to audiences in Florence, Bellingham, MIT, and Costa Rica for comments, to Fabrizio Cariani, Adam Elga, Caspar Hare, Tom Kelly, Agustin Rayo, and Miriam Schoenfield for feedback, to Sophie Horowitz for digging up this lost paper, and Jennifer Lackey for organizing the conference that got me to revive it.

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press

Pryor, J. 2000. 'The Skeptic and the Dogmatist.' *Noûs*, 34: 517–49.
Skyrms, B. 1999. *Choice and Chance: An Introduction to Inductive Logic*. Belmont, CA: Wadsworth.
Stove, D. 1973. *Probability and Hume's Inductive Skepticism*. Oxford: Clarendon.
Van Cleve, J. 1984. 'Reliability, Justification, and the Problem of Induction.' *Midwest Studies in Philosophy*, 9: 555–67.
Van Fraassen, B. 1989. *Laws and Symmetry*. Oxford: Oxford University Press.
Vogel, J. 2000. 'Reliablism Leveled.' *Journal of Philosophy*, 97: 602–23.
White, R. 2006. 'Problems for Dogmatism.' *Philosophical Studies*, 131: 525–57.
—— 2010. 'Evidential Symmetry and Mushy Credence.' In T. Szabo Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology*, pp. 161–86. Oxford: Oxford University Press.

ROGER WHITE is an Associate Professor of Philosophy at MIT whose research interests include traditional and formal epistemology.

https://doi.org/10.1017/epi.2015.9 Published online by Cambridge University Press