

INDUCTION AND THE GLUE OF THE WORLD

Harjit Bhogal*

Forthcoming in Australasian Journal of Philosophy†

Abstract

Humean metaphysical views have often been criticized for leading to inductive skepticism. If there is no glue holding the world together, as the Humean believes, then there seems to be no basis on which to infer from past to future. However, Humeans have typically been unconcerned. After all, they say, *everyone* has a problem with induction. But, if we look at the connection between induction and explanation, we can develop the problem of induction in a way that hits the Humean, but not the anti-Humean. The Humean faces an ‘internal’ problem with induction – skepticism about important inductive inferences naturally flows from their position in a way that it doesn’t for the anti-Humean. This is a major problem, perhaps a fatal one, for the Humean.

Keywords: Induction, Explanation, Humeanism, Laws of Nature

*Thanks to Thomas Blanchard, Chris Dorst, Katie Elliott, Nina Emery, Alison Fernandes, Ned Hall, Mike Hicks, Siegfried Jaag, Christian Loew, Barry Loewer, Elizabeth Miller, Zee Perry, Erica Shumener, Peter Tan and two referees for, and the editor of, *AJP* extremely helpful comments.

†This is a preprint of an article whose final and definitive form will be published in the Australasian Journal of Philosophy. The Australasian Journal of Philosophy is available online at: <http://www.tandf.co.uk/journals/>.

1 INTRODUCTION

The core idea of *Humeanism* (as that term is used in modern metaphysics and philosophy of science) is that there are no necessary connections between distinct existences. The picture that flows from this is one where the world is, at its base, just a mosaic of disconnected events pushed up against each other in spacetime [Lewis, 1986, pp. ix-xi]. There are no necessary connections ‘gluing’ the events together. In particular, laws of nature and causal relations do not provide such glue — rather, the Humean typically accepts that laws and causal relations reduce to patterns of events.

This is a popular view. However, many philosophers (most notably Armstrong [1983]) have worried that Humeanism has a problem with *induction*. If the future is disconnected from the past, in the way that Humeanism says that all events are disconnected from each other, then it’s hard to see why we should infer claims about the future from past evidence.

But most modern Humeans don’t seem especially concerned. Ned Hall, for example, calls such worries about induction ‘silly’ [2010, p. 31]. And Barry Loewer quickly rejects them as ‘question-begging’ [1996, p. 190]. The general Humean attitude seems to be that everyone faces the problem of induction, so the Humean has no special problem. This view comes out clearly in the Loewer article, but also in Beebe [2011]. She influentially argues that anti-Humeans don’t have a solution to the problem of induction — even if we accept necessary connections there is still a problem with inferring future claims from past evidence. She concludes that the Humean ‘is no more guilty than the realist [about necessary connections] of susceptibility to the problem of induction.’ (p. 526)

But I don’t think this is right – Humean and the anti-Humean are not the same footing with respect to induction. Although they are on the same footing with respect to *one version* of the problem of induction — neither can convince the inductive skeptic — I claim that

the Humean faces certain *internal problems* regarding induction that the anti-Humean does not. And these problems are very serious, perhaps fatal.

In section 2 I will set up the dialectic. In section 3 I'll discuss two types of skeptical problems that one might face — the first being the problem of convincing the skeptic that you have knowledge, and the second being the problem of building a coherent theory, attractive by your own lights, that implies you have knowledge. In sections 4-5 I'll give one argument that the Humean has difficulties with this second type of problem regarding inductive skepticism (while the anti-Humean does not). Ultimately though, that argument fails. In section 6 I'll consider an adapted argument that, I think, does suggest that Humean has a substantial problem with induction.

2 ARMSTRONG, BEEBEE, AND NECESSARY CONNECTIONS

Take *Humeanism* to be the view that there are no necessary connections between distinct existences. This is consistent with necessary connections between facts about a whole and facts about its parts, for example, since the whole and the parts are not distinct existences. But not with necessary connections between the properties of mass and acceleration since those properties are (presumably) distinct existences.

A variety of positive views are consistent with this denial of necessary connections, but, typically, modern Humeans accept that the world, at its fundamental level, is a mosaic of local events in spacetime with no necessary connections between them.

The concern for the Humean is, very roughly, that since they deny necessary connections there is nothing 'gluing' the past to the future. In fact, there is nothing 'gluing' any event to any other distinct event. And it's hard to see why we should use the past as evidence about

the future if the future is disconnected from the past. The anti-Humean, on the other hand, seems to have better prospects for dealing with induction. Since they accept the existence of necessary connections they at least have some material to work with.

The now-standard anti-Humean strategy with respect to induction – suggested by Armstrong [1983, pp. 51-60] – appeals to these necessary connections. His idea is that the best explanation of the regularities that we observe is that there are necessary connections. Exactly how this is developed depends upon the details of what necessary connections the anti-Humean accepts. But Armstrong's version says that when all the Fs we have observed are Gs the best explanation for this is that there is a *necessitation relation* that holds between F and G. (On the approach of Armstrong, Dretske [1977] and Tooley [1977] this constitutes a law that all Fs are Gs.) So, by inference to the best explanation, we are justified in believing that there is a necessitation relation between F and G. And given this, we are justified in believing that all Fs will be Gs, even those Fs we have not observed yet.

Among Humeans, however, there is a lot of doubt, perhaps even incredulity, that this anti-Humean strategy has solved the problem of induction — that this problem that we've been puzzled about for nearly 300 years is no longer an issue. In particular, Beebe [2011] argues in detail against the anti-Humean solution, but the basic point is simple — even if we accept the existence of necessitation relations, it's not clear why someone who is skeptical about induction should believe that the best explanation of the observed regularity between F and G is the hypothesis that a necessitation relation holds *eternally* between F and G, rather than some weaker hypothesis that implies that F and G have been related by necessitation *so far*.

There is still back and forth over this kind of objection to Armstrong's approach and related strategies (for example, Hildebrand [2016], Psillos [2017], Castro [2016]). But, I'm just going to assume, in the Humean's favor, that Beebe is correct about Armstrong's strategy.

However, I'll argue that, still, the Humean is in a much worse position with respect to inductive skepticism than the anti-Humean.

3 TWO APPROACHES TO THE SKEPTIC

Let's start by looking at how Beebe characterizes the problem of induction and the dialectal position she is taking. She is admirably clear about this. Her paper starts as follows:

For the purposes of this paper, I take the problem of induction to be a genuine sceptical problem. The challenge is to provide a reason to believe that inductive inferences are rational – a reason that does not beg the question against the sceptic by enshrining presuppositions that the sceptic will reject. [2011, p.504]

Beebe is understanding the problem of induction as, effectively, the problem of convincing the skeptic — of finding an argument that starts from premises that the skeptic will accept to the conclusion that induction is justified.

Traditionally, debates over skepticism were construed in this way. But, of course, convincing the skeptic is extremely hard. Nowadays discussions of skepticism don't typically aim to convince the skeptic – rather the aim is to build a coherent and attractive theory which implies that knowledge is possible, regardless of whether the skeptic is convinced. (See Nozick [1981, pp. 15-16] for an elegant description of this shift.) When the problem of skepticism is understood in this way then there is no problem with your theory 'enshrining presuppositions that the sceptic will reject'.

So we have, in effect, two skeptical problems. The first is about convincing the skeptic. The second is about building a coherent theory, attractive by your own lights, that implies that knowledge is possible regardless of whether the skeptic is convinced.

When Beebee argues against Armstrong's solution to the problem of induction she is arguing that the anti-Humean (just like the Humean) can't convince the skeptic. As I noted, I'm assuming she's right about that.

But, I'll argue that the Humean has a worse problem – they don't have a coherent theory that is attractive by their own lights and implies the right results about when inductive knowledge is possible.

The terminology I'll use for this is that the Humean has an *internal* problem with inductive skepticism – the intuition is that the Humean view itself involves commitment to premises that lead to inductive skepticism. I don't want to rest a lot of weight on this notion of 'internality' though. What counts as 'part of' or 'internal to' the Humean view is somewhat vague, and determined by sociological facts about actual Humeans. But again, the key point is simply that the Humean doesn't have a coherent theory that is attractive to Humeans themselves and implies the right results about induction.

4 A FIRST ARGUMENT

Let's start by considering an argument that the Humean is forced towards inductive skepticism:¹

(1) If we think that there is no explanation of an observed pattern then we shouldn't believe

¹The arguments I discuss in this paper are based on explanatory considerations. Sometimes the argument that Humeanism leads to inductive skepticism is developed differently, using probabilistic considerations (e.g. Foster [1982]). The idea is that most of the possible Humean mosaics consistent with our past evidence are ones where the future is chaotic, rather than continuing to be regular. So, the Humean should not expect the regularities of the past to continue.

I'm not focusing on this version of the argument because I think it's fairly easy for the Humean to respond that the argument hits anti-Humeans too – most possible anti-Humean laws lead to chaotic-looking mosaics in the future too (though see Hildebrand [2016] for disagreement). There's more to say here, but I'm leaving this probabilistic version of the argument aside.

that this pattern will continue to further, unobserved, cases.

(2) For the Humean there is no explanation of the observed regularities.

So,

(3) The Humean shouldn't think that the observed regularities will continue to further cases.

Again, the aim is to show that the Humean has an internal problem with inductive skepticism – premises that Humeans don't want to reject lead to inductive skepticism.

Perhaps you already think that this argument can't cause an internal problem because the Humean will happily reject (2). I agree. I'm not endorsing this argument, but it will be useful to explore it in some detail in order to develop a better argument.

But let's start by motivating the premises, before we look at possible Humean responses.

4.1 DEFENDING (I)

Consider this case:

Anemia: In a medical study that looked at a sample of the population you find that everyone with a certain gene also has anemia. However, given your background knowledge of the relevant gene and of anemia, you think it is unlikely that there is an explanation of the pattern, rather, you think, it is a fluke, or coincidence.

There is a powerful intuition that to the extent that we don't think that there is an explanation of this pattern — that it's merely a coincidence — then we shouldn't expect to see this correlation between the gene and anemia outside the sample.

This is a compelling thought, but we should pause to clarify it. In particular, we should be a little more careful about what it means to say that there is no explanation of the pattern.

Someone might claim that the regularity connecting the gene and anemia in our sample will always have an explanation, because for each person in the sample there is some explanation of why they have that gene, presumably one to do with how it was inherited, and there will be some explanation of why they have anemia. Conjoining all these explanations, they might claim, constitutes an explanation of the regularity we found in the sample. Or, alternatively, they might argue that the regularity will always be explained by the lower-level configuration of atoms that grounds or constitutes the regularity.

Clearly, these types of conjunctive and constitutive explanations are not the relevant type of explanations for induction. The existence of these types of explanations don't make the pattern non-coincidental [Lando, 2017, Bhogal, fort.a]. If we think that this type of conjunctive explanation is the only explanation there is of the regularity connecting anemia and the gene in our sample then we should not expect the pattern to continue outside the sample. So when I talk about explanations of a pattern or regularity I will be excluding these conjunctive or constitutive explanations. (Precisely specifying the class of explanations that aren't relevant for induction is a difficult task (see White [2005], Lando [2017], Bhogal [fort.a]), but that won't be important in the rest of the paper.)

Given these clarifications, this case suggests that to the extent that we don't think that there is an explanation of this pattern then we shouldn't expect to see this correlation outside our sample.

There are many ways an explanation of the pattern could go. Perhaps the gene causes anemia; maybe there is some common cause; perhaps the explanation is that anemia is very common in the relevant population so it's not surprising that everyone in our sample with that gene had anemia. However it goes, we need to believe that there is some explanation if we are to believe that the pattern will continue. (But obviously, we don't need to believe in

any *particular* explanation, just that there is some explanation.)

Here is another case that suggests (1) is true:

Strawson's Screen: '[Imagine that] a true randomizing device determines the colour value of each pixel on a standard 800×400 computer screen, running on a ten-times-a-second cycle – so that each pixel can take any colour value for each 1/10th second period. On the screen it appears that there is a film showing. A woman enters a house, walks over to a stove, and puts on a kettle. Life – a world, as it were – goes on in an ordered, regular fashion, exactly as regularly as in our own world. But the image is being generated by the true randomizing device. It is pure fluke that what happens on the screen appears to tell a coherent story of a regular, ordered world, rather than filling up with – or suddenly switching to – a fizz of points of colour.' [Strawson, 1989, p. 24]

There's a lot that we might take from this passage, but let's focus on a thought about induction. If we look at the screen and see the woman entering the house, walking over to the stove and putting on the kettle, and if we know the setup — that the pixels take on colors randomly — then, it seems, we should not expect the 'film' to continue. We should not expect to see the kettle boil and for the woman to make a cup of tea. Rather, we should expect the screen to suddenly switch to 'a fizz of points of colour'. We should not expect the regularities we have seen on the screen to continue.²

Again, this seems highly plausible and, in fact, Beebe herself accepts this in her [2006]. She says that 'Given what we take ourselves to know about how the computer operates, we cannot possibly make any predictions about what will happen on the screen. (Or rather, our best prediction will be that there will be some sort of jumbled mess appearing – though

²Notice, I'm not claiming that the Screen is closely analogous to the Humean mosaic (even though Strawson does, at times, claim that). One major difference is that the output of the screen is generated by a chancy process, while that's not the case for the Humean mosaic. Rather, I'm simply using the case to help motivate the link between explanation and induction.

of course we cannot say which out of the billions of possible jumbled messes will appear.)’ (p. 528)

Generalizing away from these cases, (1) seems to be an important part of scientific practice. Imagine how strange scientific practice would be if we expected patterns that are mere flukes, or coincidences, to continue. Consider, for example, spurious correlations. Between 2000 and 2009 there was an extremely high correlation between per capita cheese consumption in the USA and the number of people who died by becoming tangled in their bedsheets. In the same period the divorce rate in Maine was highly correlated with the per capita consumption of margarine. And the volume of US crude oil imports from Norway was highly correlated with the number of drivers killed in collisions with railway trains (see <http://www.tylervigen.com/spurious-correlations> for these and many other spurious correlations).

I take it that in all of these cases we think that there is no explanation for the correlation. (Again, except for the types of merely conjunctive or constitutive explanations that we are explicitly excluding.) All of these correlations are coincidences. Consequently we do not, and should not, expect these patterns to continue. To deny this would seem to run vastly counter to scientific practice and to everyday reasoning – something that the Humean does not want to do.

The Humean, I think, should, and would, accept premise (1).

4.2 DEFENDING (2)

As I noted earlier, I think most Humeans will reject (2), so I’m not going to give a full defense of it.

But there is a clear idea underlying (2): For the Humean, the mosaic of events is taken as basic and unexplained. There is nothing standing behind or outside of the mosaic making the mosaic what it is. Consequently, observed regularities, that are part of the mosaic of events, won't have an explanation either.

We will see how the Humean might reject this soon, but first we will consider a couple of other possible responses to the argument.

5 RESPONDING TO THE ARGUMENT

How will the Humean respond to this argument? I'll consider three options.

5.1 REJECT (I)

The Humean might reject (1). That is, they might say that even if we don't believe that there is any explanation of a pattern, we should, sometimes, believe that the pattern will continue to further cases.

The discussion above seems to suggest that the Humean shouldn't reject (1). As we noted, this connection between explanation and induction seems to be a central part of scientific practice. And Humeans typically want to respect scientific practice – a view which denies central parts of science would be unattractive to them.

However, Beebe [2006] does seem to reject (1).³ As we noted, in her discussion of **Strawson's Screen** she accepts that 'we cannot possibly make any predictions about what will

³I'm not totally confident of this interpretation of Beebe, but it doesn't really matter for our purposes. Perhaps the view I discuss was actually Beebe's view, or perhaps it is just an interesting view in the spirit her claims.

happen on the screen.’ So she is skeptical about making inductive inferences in the screen case. And she seems to accept that the Humean has no explanation of the observed regularities of the world — she says that ‘the continued orderliness of nature is what Strawson calls an “outrageous run of luck”’ (p. 527).

But, she claims, there is a ‘relevant epistemological difference’ between Strawson’s screen and the Humean mosaic. The difference is that, unlike in the case of the screen, ‘we take ourselves to know (fallibly, of course) that the universe is, in fact, an incredibly ordered place’ (pp. 527-528). And this belief allows the Humean to continue to infer from the past to the future.

This seems to be a rejection of (1) — accepting that the Humean has no explanation for certain patterns that we observe, but still accepting that the Humean can reasonably expect them to continue.

However, this is a somewhat puzzling move. If we really believe that there is no explanation of why the mosaic is the way that it is — if we really think that the regularities that we have seen are merely coincidences, or flukes, or outrageous runs of luck — then that seems to undermine our initial belief that the universe is an incredibly ordered place.

Consider an analogy with **Strawson’s Screen**. I might take myself to know (fallibly, of course) that images on computer screens are incredibly ordered. This seems reasonable — after all, my computer screen, and most others, typically do exhibit such order. And I might initially have this belief about Strawson’s screen. But when I come to learn that the images on the screen are formed in the random way Strawson describes then this undermines the belief I initially had in the orderly nature of the images on this screen. I should no longer be confident that the screen is an incredibly ordered place. Similarly, Beebee’s belief that regularities that we observe are flukes, or outrageous runs of luck, should undermine her

belief in the orderly nature of the world.

Clearly there's more to say here, but I don't think that appealing to this initial belief in the regularity of the world helps the Humean to reject (1).

5.1.1 SPECIAL INFORMATION

There is a more plausible way to reject (1), but it won't help the Humean avoid the conclusion.

Premise (1) says that if we think that there is no explanation of an observed pattern then we shouldn't believe that this pattern will continue to further, unobserved, cases. But this is too strong. Imagine that all the Fs you have observed have been Gs. Then you talk to an oracle who tells you that it is a fluke that the observed Fs have been Gs but, purely coincidentally, this fluke will continue. So in the future Fs will be Gs too. Clearly, you should believe that the pattern you have observed will continue, since you trust the oracle. But you also think that there is no explanation of the observed pattern. So premise (1) is false.

Here's another case: Imagine that you find out that there is some evil demon who enjoys making flukish patterns continue. All Fs you have observed have been Gs, and you think, correctly, that there is no explanation of this. But, knowing about this demon and their preferences, you rationally believe that the pattern will continue in the future because the demon will step in and adjust things so as to guarantee that future Fs will be Gs.

What these cases suggest is that sometimes we have 'special information', analogous to the oracle or our knowledge of the evil demon, that a flukish observed pattern will continue to unobserved cases, even though we typically shouldn't expect flukish patterns to continue.

So premise (1) needs to be tweaked. It should say: 'If we think that there is no explanation of an observed pattern then we shouldn't believe that this pattern will continue to further,

unobserved, cases, *unless we have some special information that the observed pattern will continue*'.

So can the Humean can avoid inductive skepticism by saying that they have such special information that the observed regularities will continue — something similar to our knowledge of the oracle or the evil demon? It's hard to see how. If a Humean observes a pattern in the mosaic — for example, that in all cases in the past force has been equal to mass times acceleration — they won't typically have access to any oracle to guarantee that this pattern will continue. And neither do they think that there are any powerful forces, whether of an evil demon kind or otherwise, that will make force equal to mass times acceleration in the future.

So the special information clause doesn't seem to help the Humean avoid inductive skepticism.

5.2 THE ANTI-HUMEAN HAS THE PROBLEM TOO

The Humean might respond by accepting that they have an internal problem with induction but claiming that the anti-Humean does too.

That is, they might accept that (1)-(3) shows that there isn't a coherent position that is attractive to the Humean and avoids inductive skepticism. But they might claim that, similarly, there isn't a coherent position that is attractive by anti-Humean lights and avoids inductive skepticism.

The anti-Humean, though, can easily avoid the argument (1)-(3) since they think that the necessary connections they postulate provide an explanation of the observed regularities.

But aren't there still concerns for the anti-Humean here? For example, how do they know

that the laws won't change? If we are worried that the laws will change in the future then perhaps we can't be confident that the next F we see will be a G. And how does the anti-Humean know that the laws are not extremely complicated? If the laws can be so complicated then maybe don't have reason to think that all observed Fs are Gs is explained by a law that all Fs are Gs. Perhaps it is explained by some much more complicated law and this more complicated law implies that at some point in the future Fs will no longer be Gs.

These are legitimate concerns for the anti-Humean. But I don't think they are *internal problems*. That is to say, even in light of these issues, it looks like there can be coherent positions that are attractive by the lights of the anti-Humean and avoid inductive skepticism, even if others will not be convinced.

For example, consider again Armstrong's view that the best explanation of observed regularities, like all observed Fs being Gs, is that there is an unchanging law that all Fs are Gs. And so we should believe that future Fs will be Gs.

Beebe thinks that the skeptic will deny this claim about bestness of explanation – because the skeptic is worried about the possibility of laws changing or being extremely complex. But, regardless of whether Armstrong will convince the skeptic, it seems perfectly coherent for him to claim that the best explanation of all observed Fs being Gs is, in fact, the simple, unchanging, law that all Fs are Gs, rather than any law which changes over time or which is very complicated.

And further, this Armstrongian view seems to be attractive by the anti-Humean's own lights. At least, there is no obvious reason why the anti-Humean should be unhappy accepting this claim about inference to the best explanation.

So, I think Armstrong's view is an example of how the anti-Humean can have a coherent view, attractive by their own lights, that avoids inductive skepticism – regardless of whether

the Humean or the inductive skeptic will agree.

Of course, I've not ruled out the possibility that someone in the future will come up with an argument showing that anti-Humean positions face internal problems too. Perhaps, for example, someone could argue that scientific practice seriously countenances laws changing and hence the claim that the best explanation of all observed Fs are Gs is an unchanging law is in conflict with the practice. In fact, a reviewer notes that Dirac [1937] countenances the possibility of the fundamental physical constants changing. However, such suggestions are rather rare in scientific practice, and so don't seem to undermine the anti-Humean's claim about inference to the best explanation.

If an argument could be made that the anti-Humean doesn't have a coherent view that is attractive by their own lights then that would validate the response that the Humean has no special problem. But as it stands it's hard to see how such an argument could be developed.

5.3 REJECT (2)

Now let's consider perhaps the most plausible Humean response to the argument — rejecting premise (2).

The idea motivating (2) is that, for the Humean, the mosaic of events is taken as basic. There is nothing standing behind or outside the mosaic making it what it is. Consequently, observed regularities, that are part of the mosaic of events, won't have an explanation either. But the Humean will likely reject this reasoning. Although they take the mosaic to be basic they don't take every event in the mosaic to be unexplained – they don't reject all instances of scientific explanation.

How is this possible? The idea is that just because the mosaic is basic doesn't mean that

parts of the mosaic are unexplained — bigger patterns in the mosaic can explain smaller ones. In particular, Humeans typically accept that the laws of nature are general patterns in the mosaic — ones that are simple and informative⁴ (though saying exactly what simplicity and informativeness come to is very complicated) — and that these general patterns can explain events in the mosaic.

Broadly speaking then, we can think of such Humean explanation as having a *pattern subsumptionist* or *unificationist* flavor — the general patterns in the mosaic explain the particular events. Particular events are explained by fitting them into the more general patterns. This conception of explanation has a long tradition. For example, Hempel [1966, p. 488] claims that ‘The understanding [an explanation] conveys lies...in the insight that the explanandum fits into, or can be subsumed under, a system of uniformities represented by empirical laws or theoretical principles’. Kneale [1949], Friedman [1974], Kitcher [1981] and many others express similar ideas.

Loewer [1996, p. 113] explicitly accepts that Humean laws explain in this pattern subsumptionist way. As do other Humeans like Smart [2013], Miller [2015, section 4] and Bhogal [fort.b, section 2.1].

So, the Humean can accept that particular events in the mosaic are explained, even when the mosaic as a whole is not explained by anything else.

And the Humean can say that the observed *regularities* are explained in much the same way — by being subsumed to more general regularities. Consequently, the Humean will happily reject (2) — the argument (1)-(3) doesn’t cause them an internal problem.

⁴See Lewis [1986, p. 42-3], and many refinements of this idea, e.g. Cohen and Callender [2009], Loewer [1996], Hall [2010], Hicks [2018], Dorst [2019], Jaag and Loew [fort.].

6 AN IMPROVED ARGUMENT

So, the first argument doesn't cause a problem for the Humean. In this section I'm going to discuss an adapted argument that does better. This argument doesn't aim to show that Humeanism leads to throughgoing inductive skepticism though. Rather, it aims to show that the Humeanism leads to skepticism about some important inductive inferences.

One piece of terminology first: Take an observed pattern. For example, that all observed Fs are Gs. Then call a pattern that includes all the observed cases but also continues to unobserved cases an *extended pattern*. For example, the pattern that all Fs are Gs extends the pattern that all observed Fs are Gs.

Two quick clarifications about this notion. Firstly, there is a sense in which all observed Fs are Gs extends to larger pattern *whatever comes about*. For example, there is a sense in which the pattern 'Every F was G until May of 2030, then every F was H instead' counts as an extension of 'all observed Fs are Gs'.⁵

But there is a different, perhaps more everyday, sense of what it is for a pattern to extend, or continue, where all Fs are Gs counts as an extended pattern of all observed Fs are Gs but 'every F was G until May of 2030, then every F was H' doesn't count as an extended pattern. It is this second sense of 'extends' that I'm using.

Secondly, the universal generalization all Fs are Gs is intuitively an extension of the pattern that 20 Fs we have observed have been Gs. However, the universal generalization doesn't consist in merely adding additional instances to the smaller pattern – universal generalizations are not merely conjunctions of instances. In addition they require a 'that's all' clause specifying the lack of any cases that violate the generalization. It's useful to keep this in

⁵Thanks to a reviewer for this point.

mind when evaluating the premises to come. (In particular, see footnote 10.)⁶

Given this, we can formulate the argument:

(4) If we think that there would be no explanation of the extended pattern (even if it were to hold) then we shouldn't think that the observed pattern will continue to the extended pattern, unless we have some 'special information' that the observed pattern will continue.

(5) The Humean thinks that there would be no explanation of the most general regularities of the world, no matter what they end up being.

(6) The Humean has no 'special information' that the observed pattern will continue.

So,

(7) The Humean shouldn't think that the observed regularities will extend to the most general regularities of the world.⁷

This argument is more complicated than (1)-(3), so the premises need some clarification, as well as defense. But here's the rough guiding intuition: For the Humean nothing explains the most general regularities of the world. And if nothing explains those regularities then we, who have only observed some instances of those regularities, shouldn't expect those general regularities to come about.

6.1 PREMISE (4)

The formulation of premise (4) is a little complicated, but here's the motivating idea: For an explanation to license your belief that a pattern will extend from observed cases to unobserved cases that explanation must itself extend, so to speak, to those unobserved cases.

For example, consider again **Anemia** and the correlation between anemia and the gene in

⁶Thanks to a reviewer for discussion of this point.

⁷Thanks to a reviewer for suggesting formulations of some of these premises.

our sample. (1) says that to believe that the pattern will continue outside the sample we must think that there is some explanation of the correlation in the sample. This seems right (modulo the discussion of ‘special information’ in section 5.1.1). But we can say something stronger.

If we believe that the correlation will continue outside the sample, then we must think that there is some explanation of the correlation *that applies outside the sample*. For example, imagine an adapted **Anemia** case, where there is an explanation of the correlation between anemia and the gene that we observed, but it’s just that some bias in the way the sample was formed led to the correlation. If we think that this is the only explanation of the correlation then we shouldn’t believe that the correlation will continue outside the sample. We shouldn’t believe that the pattern will continue to unobserved cases because we don’t think that there is an explanation that extends to unobserved cases.

Another example: Imagine a variant of Strawson’s screen. This screen has been designed to show a movie, played from a standard computer, for the first 5 minutes after it’s turned on. But after this time each pixel takes on a color value at random as Strawson described. If we, knowing this, watch the film for the first 2 minutes then we should believe that the patterns of the movie will continue for 3 more minutes, but after that we shouldn’t expect the patterns to continue. Our beliefs about how a pattern will continue to unobserved cases should mirror our beliefs about how far an explanation will extend.

That’s the motivating idea. How do we formulate it precisely? Well, what is it to think that there is no explanation of an observed pattern that extends to unobserved cases? It is to think that, even if the unobserved cases did fit with the observed pattern then that would be merely coincidental — there would be no explanation of the extended pattern.⁸ And so

⁸We have to be a little careful here, because there are two ways of reading this ‘even if’ conditional. Take

premise 4 (ignoring for a moment the clause about ‘special information’), says that in such a case – a case where we think that there would be no explanation of the extended pattern, even if it were to come about – we should not expect the observed pattern to continue.

This way of formulating the idea is very much in the spirit of Roger White’s [2005] discussion. For example, he discusses observing a run of heads coin tosses: ‘Suppose we are quite certain that the coin is evenly weighted and tossed in the ordinary way, and so if it does land heads every time this is just a fluke. In this case we have no reason to expect any other tosses [other than the ones we have observed] to land heads.’ (p. 10) That is to say, if there would be no explanation of the extended pattern of coin tosses landing heads, then we shouldn’t expect the observed pattern of coin tosses landing heads to continue.⁹

So, the formulation of premise (4) expresses this idea that if we don’t believe that there is an explanation of the observed pattern that extends to the extended pattern then we shouldn’t believe that an observed pattern will continue to the extended pattern.¹⁰

So far we’ve ignored the clause ‘unless we have some special information’. That clause is included for exactly the same reasons as discussed in section 5.1.1.

Now we can see clearly what (4) comes to it’s hard to see how a Humean could be comfortable rejecting it. (4) seems to be a part of scientific practice that the Humean should, and would, accept.

the anemia case again. We think that the correlation between the gene and anemia is a coincidence. It seems natural, then, to say that if the correlation extended outside the sample that would be a coincidence. But it’s also natural to say that if the correlation extended outside the sample then there must, contrary our initial claim, be some explanation for it. I mean the conditional to be read in this first way, where we where we hold fixed our initial beliefs about the explanatory structure of the situation.

⁹Peacocke [2003, Chapter 5] expresses a similar idea.

¹⁰ And this idea, I take it, is just as plausible when the extended pattern is a universal generalization and so involves a ‘that’s all’ clause.’ Consider, for example, a version of the Anemia case where the extended pattern under consideration is a universal generalization, say the generalization ‘everyone in my city with that gene has anemia’. I take it that if we think that there would be no explanation of this extended pattern, then we shouldn’t think that the pattern we observed in the sample will extend to this universal generalization.

6.2 PREMISE (5)

As we noted in section 5.3, the Humean can make sense of parts of the mosaic being explained in a broadly pattern subsumptionist or unificationist way. Particular events are subsumed to the general patterns of the mosaic. This appears to be how Humean laws explain.

But there is a limit to how far this pattern subsumption could go. Particular Fs that are Gs can, perhaps, be explained by being subsumed to the general pattern that all Fs are Gs. Similarly with certain patterns, like all observed Fs being Gs. But the general pattern itself — the fact that all Fs are Gs, obviously cannot be explained in this way. We cannot explain a pattern by subsuming it with itself.

We need to be a little careful here, because some higher-level patterns might be explained, not by subsuming them to more general patterns at the same level, but by unifying them with lower-level patterns. For example, take some chemical regularity — all reactions of magnesium with oxygen are exothermic, for example. Physical chemists can subsume this fact to regularities at the physical level.

But again, this process of subsumption to more general patterns has to end somewhere. At some point we are left with the *most general regularities of the world* — those cannot be explained by subsuming them to more general patterns. And, given the Humean picture, those most general patterns can't be explained in any other way either. There is nothing standing outside the mosaic to explain those patterns. (As we initially discussed in 4.1 we are still excluding constitutive explanations, or those that purport to explain a regularity by merely conjoining explanations of the instances. The premises should be understood with this in mind.)

So, it looks like the Humean will have to accept (5) — they will say that there would be no

explanation of the most general regularities of the world, whatever they end up being.

6.3 PREMISE (6)

The notion of special information at work here is exactly the same as the one discussed in section 5.1.1. As we noted, it doesn't look like the Humean has any information, analogous to that of the oracle or the evil demon, that the observed pattern will continue.

6.4 (7) AND A PROBLEM WITH INDUCTION

But putting (4), (5) and (6) together gets you (7). And (7) is a problem for the Humean.

Imagine that all Fs are Gs is one of those most general regularities of the world referred to in (5). Then it seems that the Humean shouldn't think that all observed Fs are Gs will continue to the extended pattern that all Fs are Gs, because there is no explanation of this extended pattern.

So, it seems that the Humean should be skeptical of at least a certain type of inductive inference – inferences from observed regularities to the most general regularities of the world. For example, perhaps it is one of the most general regularities of the world that in all interactions energy is conserved. If so then the Humean should be skeptical of inferring from the fact that in all observed interactions energy has been conserved to in all interactions energy is conserved.

Again, this is an internal problem for the Humean. The issue is not that the Humean cannot convince an intransigent skeptic. It's that principles that they themselves want to accept lead them to be skeptical of certain inductive inferences. And there does not seem to be analogous problems for the anti-Humean. As we noted in section 5.2, the anti-Humean

has problems convincing the skeptic, but they seem to have a coherent view, attractive by their own lights, that avoids inductive skepticism.

It's important to be clear here about what the problem is for the Humean. The argument here doesn't raise problems for every instance of induction for the Humean. The argument works by claiming that we cannot infer from an observed pattern to an extended pattern where we don't think that there's an explanation of the extended pattern. But for many inductive inferences Humeans can claim that there is an explanation of the extended pattern. For example, imagine that there are 10,000 ravens in the history of the world and they are all black (although you don't know this). You've seen 20 ravens and they are all black, so you infer that the next 5 ravens you see will be black. The argument I've given doesn't cause a problem for this inference, because here the extended pattern is that 25 ravens are black, and for the Humean who accepts pattern subsumptionism that pattern could be explained by subsuming it to the pattern that all ravens are black.

In light of this, perhaps the Humean might accept the argument (4)-(7) but declare themselves unconcerned — sure they have problems with some very large scale inductive inferences, but there is no problem for most piecemeal inferences. But I don't think this is a particularly satisfying response. It is a problem if the Humean cannot infer from the fact that in all observed interactions energy has been conserved to fact that in all interactions energy is conserved. More generally, it's a problem if a scientist can't use past observations to become justified in believing very general regularities about the world. These inductive inferences to the most general patterns may be rare in everyday life but they are vitally important when we are forming beliefs about the basic structure of the world. To be skeptical about these inferences would undermine our fundamental science.

Given the argument (4)-(7), then, it looks like Humeans really do face a problem with

induction. The objection that they should be skeptical of certain inductive inferences isn't a silly one, and it can't easily be brushed off as question-begging. If the Humean cannot develop an adequate response then we have a powerful reason to reject Humeanism.

7 CONCLUSION

The Humean, I've argued, faces a problem with induction. And it's not merely the problem that faces everyone – that they cannot convince an intransigent skeptic. Rather it's that premises they want to accept lead them to be skeptical about important inductive inferences. This is a big problem for the Humean, perhaps a fatal one.

It's the last paragraph, so time to come clean: I am a Humean, and this whole paper has been about trying to understand the best version of this objection to my view. So what's my favored response to the argument? My hope is that a partial rejection of (4) is defensible. The thought being that (4) – the connection between explanation and induction – applies in most cases, but the Humean mosaic is a special case. There is a way in which the most general regularities of the world, for the Humean, have a kind of special status such that (4) doesn't apply. Clearly, though a very substantial case would have to be made to defend the Humean here, and that's a task for another time.¹¹

References

David Armstrong. *What is a Law of Nature?* CUP, 1983.

¹¹In particular, it's a task for my manuscript 'Does anything explain the regularity of the world?'

- Helen Beebe. Does Anything Hold the Universe Together? *Synthese*, 149(3):509–533, 2006.
- Helen Beebe. Necessary Connections and the Problem of Induction. *Nous*, 45(3):504–527, 2011.
- Harjit Bhogal. Coincidences and the Grain of Explanation. *Philosophical and Phenomenological Research*, fort.a.
- Harjit Bhogal. Nomothetic Explanation and Humeanism About Laws of Nature. *Oxford Studies in Metaphysics*, fort.b.
- Eduardo Castro. Is the Humean Defeated by Induction? A Reply to Smart. *Philosophia*, 44(2):435–446, Jun 2016.
- J Cohen and C Callender. A better best system account of lawhood. *Philosophical Studies*, 2009.
- Paul AM Dirac. The cosmological constants. *Nature*, 139(3512):323–323, 1937.
- Chris Dorst. Towards a Best Predictive System Account of Laws of Nature. *The British Journal for the Philosophy of Science*, 70(3):877–900, 2019.
- Fred Dretske. Laws of nature. *Philosophy of Science*, 44(2):248–268, 1977.
- John Foster. Induction, explanation, and natural necessity. *Proceedings of the Aristotelian Society*, 83(n/a):87–101, 1982.
- Michael Friedman. Explanation and Scientific Understanding. *The Journal of Philosophy*, 71(1):5–19, 1974.
- Ned Hall. Humean Reductionism About Laws of Nature. 2010.

- Carl G. Hempel. *Philosophy of Natural Science*. Englewood Cliffs, N.J., Prentice-Hall, 1966.
- Michael Townsen Hicks. Dynamic humeanism. *British Journal for the Philosophy of Science*, 69(4):983–1007, 2018.
- Tyler Hildebrand. Natural properties, necessary connections, and the problem of induction. *Philosophy and Phenomenological Research*, pages 668–689, 2016.
- Siegfried Jaag and Christian Loew. Making best systems best for us. *Synthese*, pages 1–26, fort.
- Philip Kitcher. Explanatory unification. *Philosophy of Science*, 48(4):507–531, 1981.
- William Kneale. *Probability and Induction*. OUP, 1949.
- Tamar Lando. Coincidence and common cause. *Noûs*, 51(1):132–151, 2017.
- David Lewis. *Philosophical Papers, Volume 2*. Philosophical Papers, Volume 2. Oxford University Press, Oxford, 1986.
- Barry Loewer. Humean supervenience. *Philosophical Topics*, 24(1):101–127, 1996.
- Elizabeth Miller. Humean Scientific Explanation. *Philosophical Studies*, 172(5):1311–1332, 2015.
- Robert Nozick. *Philosophical Explanations*. Philosophical Explanations. Belknap Press, Cambridge, MA, 1981.
- Christopher Peacocke. *The Realm of Reason*. Oxford University Press, 2003.
- Stathis Psillos. Induction and natural necessities. *Journal for General Philosophy of Science*, 48(3):327–340, 2017.

Benjamin T. H. Smart. Is the humean defeated by induction? *Philosophical Studies*, 162(2): 319–332, 2013.

Galen Strawson. *The Secret Connexion: Causation, Realism, and David Hume*. Oxford University Press, 1989.

Michael Tooley. The Nature of Laws. *Canadian Journal of Philosophy*, 7(4):667–698, 1977.

Roger White. Explanation as a guide to induction. *Philosophers' Imprint*, 5:1–29, 2005.