

NECESSARY CONNECTIONS AND THE PROBLEM OF INDUCTION¹

Helen Beebe

Please do not cite this version. The published version is: ‘Necessary Connections and the Problem of Induction’, *Nous*, 45 (2011): 504-27.

1. Introduction

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. The proposal that I shall consider and reject in this paper is that the sceptical problem can be solved if we focus on the metaphysics of laws of nature.

Why might one think that the problem of induction has a metaphysical solution? Well, one thing that might be worrying the inductive sceptic is this: here we are, cosily occupying a tiny corner of the vast reaches of Time. Everything’s been going along pretty nicely up to now – but *it might all fall apart! Anything could happen!* After all, what’s stopping it? Stuff happens. So far, stuff has been kind enough to happen in nice, regular, predictable ways, by and large. But maybe the regularity of the Universe thus far has just been a matter of cosmic luck, and maybe next year or next week or in the next ten minutes our luck will run out and chaos will descend – or maybe the Universe will start behaving in other regular but far less friendly ways. Simon Blackburn calls this unfortunate condition ‘inductive vertigo’ (1993, 98). What the vertigo-sufferer apparently needs is a metaphysician, for only a metaphysician is in a position to tell the afflicted that, in fact, it *can’t* all fall apart.

It looks as though the kind of metaphysician who is in a position to offer a cure for inductive vertigo will be someone who holds that there is something in the world that *makes* it regular: something that constrains how things can happen in such a way that they are guaranteed not to fall apart. In other words, it seems that the vertigo-sufferer's best bet is to consult a necessitarian of some sort. And here she has a variety of options, of whom the two most prominent are:

(a) David Armstrong. On Armstrong's view, its being a law that *F*s are *G*s is a matter of the universals *F* and *G* being related by a higher-order universal, *N* ('*N*' for 'necessity'). Their being so related is supposed to *guarantee* that all *F*s are *G*s. So, assuming that the world is a world of laws, it is guaranteed to be a regular world. (See Armstrong 1983; also Tooley 1977 and Dretske 1977.)

(b) Brian Ellis. Ellis calls his view 'scientific essentialism'. On this view, to be a member of a natural kind (an electron, a water molecule, a carbon atom) is to be intrinsically and essentially disposed to behave in certain kinds of ways in certain circumstances. Electrons are, by their very nature, disposed to repel each other. So any possible world that contains electrons will be a world in which electrons repel each other: the nature of an electron guarantees that it will behave the same way given the same conditions. (See Ellis 2001 and 2002.)

Armstrong and Ellis have both argued that the necessitarian can, while the Humean cannot, solve the problem of induction.ⁱⁱ Given what I said above, it is easy to see why this is a tempting thought, since it looks as though inductive vertigo is a peculiarly Humean affliction. Humeans – by which I mean those philosophers who

refuse to allow necessary connections into their ontology – can offer no metaphysical glue to stop things falling apart. On a Humean view, *nothing* stops things falling apart: the regularity of the Universe is a brute, inexplicable fact.

Necessitarians, whether of an Armstrongian or an Ellisian variety, make two claims. First, they make a claim about what it is to be a law of nature. (For Armstrong, laws are contingent relations of necessity holding between universals. For Ellis, they are facts about the essential dispositions of natural kinds.) Second, they make the additional claim that there are, in fact, laws of nature, so characterised. In the context of the problem of induction, it is the question of whether there is any sceptic-busting justification for believing the second claim that is of interest. If you genuinely believe that there are timeless necessary connections, or that the world is composed of natural kinds with unchangeable essences – that is, if you believe both of the above claims – you aren't going to suffer from inductive vertigo. But, in the context of the problem of induction, that's not terribly interesting. The pertinent question is whether necessitarianism offers a *cure* for inductive vertigo. That is, can a vertigo-sufferer be *persuaded* to believe in the kinds of necessary connection whose obtaining guarantees that things won't fall apart?

The purpose of this paper is to argue that the answer to this question is no: if there is a problem of induction for Humeans, there is also a problem for necessitarians. So, as far as the problem of induction is concerned, realism about necessary connections does not have the advantage that Armstrong and Ellis claim for it.

2. Explaining regularity

The central feature of Armstrong's necessitarian account of inductive inference, as we

shall see, is that it invokes inference to the best explanation (IBE). How might we invoke IBE as a way of getting us to the existence of necessary connections? Well, the Universe is an extraordinarily regular place. It is this amazing regularity that the existence of necessary connections is alleged to explain.ⁱⁱⁱ

One way to motivate this claim is the idea that to explain why *P* is the case is to show that *P must* be the case.^{iv} This is a mode of explanation that we appeal to sometimes in ordinary life: to the question, ‘why did you do that?’, someone might reply that they didn’t really have a choice: they *had* to do it. And to the extent that we buy their story about not having a choice, we count this as a sufficient explanation of why they acted in the way they did. More generally, if we think of the question, ‘why *P*?’ as a request for an explanation of *why P rather than not-P*, then we can see why ‘because *P has* to be the case’ counts as an explanation: if we know that not-*P* is not a genuine possibility, then, in some sense at least, we know why *P*.

In the particular case of explaining the regularity of the Universe, then, one candidate explanation is that the Universe has a nature such that it *must* be regular: given the underlying nature of things, it could not be anything other than regular. And to say that given the underlying nature of the Universe, it could not be anything other than regular is – very broadly – to say that there are necessary connections in the Universe.

Of course, we haven’t yet established that the existence of necessary connections is the *best* explanation for the regularity of the Universe; only that it is a *candidate* explanation. The further claim that it’s the *best* explanation is motivated by the thought that there is simply no other candidate explanation available. Humeans hold – implicitly at least – that there is no explanation for why the Universe is regular: it’s just a brute fact that things happen in nice, predictable ways. So of course if we

have a choice of only one possible explanation, then that explanation is, by default, the best explanation. QED.

It's important that we distinguish at this point between explaining the *general* regularity of the Universe – explaining why things in general keep on ticking along rather than falling apart – from explaining why some *particular* regularity obtains. 'Why is the Universe regular?' is a different question to 'why are all *F*s *G*s?'. And this opens up the possibility that even if the Humean can't provide an answer to the first question, she can perfectly well provide an answer to the second. Since the kinds of arguments for justifying induction I'm going to be considering rely on the thought that necessary connections explain particular regularities, we need to see whether Humeans *can* provide an answer to the second question – since if they can, the necessitarian's explanation for the obtaining of particular regularities will not be the only explanation, and hence will not simply be the best explanation by default.

I take it that the Humean is, in general, perfectly capable of availing herself of the kinds of explanation of regularities that are ordinarily given in everyday and scientific contexts. For example, if I ask why emus don't fly, and you reply that they don't fly because they can't, I'm prepared to accept that you've done something explanatory; but on the other hand, you certainly haven't given me a *decent*, let alone the best, explanation of why emus don't fly. A good explanation would tell me something about aerodynamics and about what kind of wing a bird with an emu's size and weight would need to have in order to get off the ground. In other words, most explanations of regularities consist in fitting the regularity in question into some deeper or more general regularity, or of telling a story about the mechanisms via which *F*-ness gets to cause *G*-ness. Such explanations need not appeal to any necessary connections.

Are there *any* cases where a necessary connection between particular properties *F* and *G* might count as the only possible, and hence the best, explanation of the fact that all *F*s are *G*s? The only possible cases, it seems to me, are cases where ‘all *F*s are *G*s’ has the status of a fundamental regularity: something not reducible to or subsumable under some deeper or more general regularity. In such a case, as with regularity in general, it might be argued that since the necessitarian can give an explanation of sorts for why the regularity obtains, whereas the Humean can’t, the necessitarian explanation gets to count as the best explanation by default.

Of course, Humeans might want to deny that necessary connections explain the existence of regularities. For example, one might – drawing inspiration from a traditional interpretation of Hume – want to deny the very intelligibility of the notion of necessary connection, in which case the postulation of such things is no explanation at all of anything.^v One might also be sceptical about IBE either in general or in the specific context of the justification of induction, as is Bas van Fraassen.^{vi}

For the purposes of this paper, however, I want to grant to the necessitarians that the existence of necessary connections is the best explanation for the regularity of the Universe, and also for the obtaining of particular, fundamental regularities, and that this provides a reason to believe in them – provided, of course, that we have a reason to believe that such fundamental regularities exist in the first place. I shall argue that, even so, necessitarians do not have a distinctive way of solving the problem of induction available to them. In other words, even if necessary connections are indeed the best explanation for the regularity of the Universe, still the necessitarian solution to the problem is no good.

3. Armstrong's solution to the problem of induction

In his *What is a Law of Nature* (1983), Armstrong argues that his brand of realism about necessary connections can, while a Humean account of laws cannot, solve the problem of induction. '[I]f laws of nature are nothing but Humean uniformities', he says, 'then inductive scepticism is inevitable' (1983, 52), whereas 'the Universals theory can do better' (1983, 104). I'll first quickly outline Armstrong's view of laws, then, second, sketch his argument about induction, and, finally, criticise that argument.

Armstrong's view of laws

What is the difference between its being a *law* that all *F*s are *G*s and its merely being the *case* that all *F*s are *G*s? Armstrong's answer runs as follows. Its being a law that all *F*s are *G*s – *F* and *G* are universals here – is a matter of there being a second-order universal *N* ('*N*' for necessity) that relates the first-order universals *F* and *G*.

(Armstrong writes this '*N(F,G)*': *F*-ness necessitates *G*-ness.) So whenever we have an instance of *F*, it is guaranteed, by *N*, that we will also have an instance of *G*. By contrast, its merely being the case that all *F*s are *G*s just requires that instances of *F* are always, in fact, accompanied by instances of *G*: there is no *necessary* connection between the two.

On a regularity – that is, Humean – view of laws, by contrast, its being a law that all *F*s are *G*s is either merely a matter of its being the case that all *F*s are *G*s (this is the 'naïve regularity theory'), or else it is a matter of 'all *F*s are *G*s' having some special status: being an axiom or theorem in the best systematisation of what happens in the Universe, for example.

Armstrong's solution to the problem of induction

For simplicity, let's just stick with the contrast between Armstrong's view and the naïve regularity theory. Here's how Armstrong's argument goes. Suppose we thought that inductive inference was just a one-step inference from 'all observed *F*s have been *G*s' to 'all *F*s are *G*s'. Armstrong thinks that if *that's* what inductive inference amounts to, then it cannot be justified. However, he claims that we can conceive of inductive inference differently: as a two-stage inference. The first stage is inference to the best explanation. From 'all observed *F*s have been *G*s', we infer, via IBE, that it is a law that all *F*s are *G*s. Now – and this is the second step – its being a law that all *F*s are *G*s entails that all unobserved *F*s are *G*s, since the unobserved *F*s are just a subset of the *F*s. So IBE and straightforward entailment together deliver the conclusion, 'all unobserved *F*s are *G*s', from the premise that all observed *F*s are *G*s.

Deductive inference is paradigmatically rational, so there's no problem with the second stage of the argument.^{vii} So the rationality of inductive inference hangs on whether the first stage of the argument is rational. Armstrong argues that the first stage *is* rational if we believe in necessary connections, but it isn't if we are naïve regularity theorists. And the reason is that the inference from 'all observed *F*s are *G*s' to 'it is a law that all *F*s are *G*s' only gets to be a case of genuine IBE if we assume realism about *N*.

The inference isn't genuine IBE for the naïve regularity theorist, Armstrong claims, because for the naïve regularity theorist, 'it is a law that all *F*s are *G*s' simply means 'all *F*s are *G*s', which in turn is logically equivalent to 'all observed *F*s are *G*s and all unobserved *F*s are *G*s'. Given this, the naïve regularity theorist who attempted to invoke IBE in the move from 'all observed *F*s are *G*s' to 'it is a law that all *F*s are *G*s' would in effect be claiming that 'all observed *F*s are *G*s' is explained by 'all

observed *F*s are *G*s and all unobserved *F*s are *G*s'. But neither conjunct of the alleged explanans really explains the explanandum. The first conjunct – 'all observed *F*s are *G*s' – just *is* the explanandum, and nothing explains itself. And the second conjunct – 'all unobserved *F*s are *G*s' – manifestly doesn't explain why all observed *F*s are *G*s. Given this, the conjunction of the two doesn't explain the explanandum either; so inference from 'all observed *F*s are *G*s' to 'it is a law that all *F*s are *G*s' is not, for a naïve regularity theorist, an instance of IBE. So (in the absence of an alternative, Humean-friendly solution to the problem of induction, which would of course obviate the need for the two-stage model in the first place) our inference from the observed regularity to the law is not rational. Hence the naïve regularity theorist can't solve the problem of induction by adopting the two-stage model of inductive inference.

For the necessitarian, by contrast, the inference from 'all observed *F*s are *G*s' to 'it is a law that all *F*s are *G*s' *is*, Armstrong claims, an instance of IBE. That *F* and *G* are necessarily connected really does explain – indeed is the best explanation for – why all observed *F*s have been *G*s. So the problem of induction *can* be solved by adopting the two-stage model of inductive inference.

What's wrong with the argument

I granted earlier, for the sake of the argument, that inference from the existence of fundamental regularities to the existence of necessary connections is indeed an instance of IBE and hence (again for the sake of the argument) rational. Now, the regularity that needed explaining in that case was regularity *simpliciter* – regularity across all of space and time – and the necessary connections whose existence were posited in the explanans were (implicitly) *timeless* necessary connections: necessary connections that hold across all of space and time. But in the context of inductive

inference our *explanandum* is not regularity *simpliciter*; what calls for explanation is not that all *F*s are *G*s, but that all so-far *observed F*s have been *G*s. For of course, prior to a satisfactory solution to the problem of induction, the fact that *all F*s are *G*s is not yet something that calls for explanation, since we do not yet have any reason to suppose that it is a fact.

The fact that what calls for explanation is only that the observed *F*s have been *G*s is important, since alternative explanations come into play, aside from the one that postulates timeless necessary connections. In particular, consider the following alternative explanation ('SF' for 'so far'):

(SF) *F* and *G* have been necessarily connected so far,

which contrasts with Armstrong's proposed explanation ('T' for 'timeless'):

(T) *F* and *G* are timelessly (eternally) necessarily connected.

Of course, one might immediately object that (SF) entails (T), since any necessary connection that has held so far is guaranteed, in virtue of being a *necessary* connection, to hold for all times. I address this objection in more detail in §4 below, but for now, let us grant that it is possible for (SF) to be true and (T) false. For example, we might postulate necessary connections that will exist only until next Tuesday, or that will start relating completely different universals at 7 o'clock this evening. Each of these hypotheses renders (SF) true but (T) false. So (pending the argument of §4) (SF) is an alternative explanation of our observed regularity.

This is not, of course, to say that (T) is no longer in the running as a candidate

explanation: if the existence of timeless necessary connections explains regularity *simpliciter*, then, plausibly, it also explains observed regularity, since regularity *simpliciter* is just observed regularity plus unobserved regularity. So far, so good. But it does not follow that (T) is the *best* explanation of observed regularity. The question, then, is whether (SF) is an *equally good* explanation for why the observed *F*s have been *G*s. If it is, then (T) will not be the *best* explanation of our observed regularity, and hence the conclusion that it is a timeless Armstrongian law that all *F*s are *G*s will not be licensed by IBE. This is important, of course, because (SF) does not satisfy the second step of Armstrong's proposed two-step inference (again, pending the argument of §4): (SF) does not entail that *all F*s are *G*s. Indeed, it doesn't even entail that the *next F* will be a *G*.

I claim that there is no reason to think that (T) explains our observed regularity any better than (SF). After all, in general, if *E* is the best explanation of *A&B*, it doesn't follow that *E* is also the best explanation of *A*. So just because $N(F, G)$ is the best explanation of the fact that *all F*s are *G*s – both observed and unobserved – it doesn't follow that it is the best explanation of the fact that all *observed F*s are *G*s. If *all* we are trying to explain is the fact that the observed *F*s have been *G*s, then the hypothesis that *F* and *G* have been necessarily connected *so far* is surely just as good a candidate explanation as is the hypothesis that *F* and *G* are timelessly necessarily connected.

If that is right, then the postulation of timeless necessary connections is not sanctioned by IBE, and Armstrong's proposed solution to the problem of induction fails. Indeed, it fails precisely because it presupposes an illicit inductive step. If IBE sanctions only inference to (SF), then an extra step, between Armstrong's first and second steps, is needed to get us from (SF) to (T). And only inductive inference can

be used to take this step.

4. Objections met

There are two broad ways in which one might attempt to save Armstrong's solution. First, one might argue that the alleged rival candidate explanation, (SF), is not a genuine rival candidate explanation at all, because there is something incoherent about the notion of a time-limited necessary connection. So the only way (SF) could be true would be for (T) to be true. Second, one might argue that, while the notion of a time-limited necessary connection makes sense, (T) nonetheless constitutes a *better* explanation of observed regularity than does (SF). I shall consider these possibilities in turn.

Is the notion of time-limited necessity coherent?

First, then, the defender of Armstrong might attempt to claim that there is something incoherent about the notion of necessity that is both genuine *necessity* and yet also time-limited in some way. After all, one might protest, it's surely in the *nature* of natural necessity that it is timeless. At this point, we need to examine in a little more detail the precise options that are available if we are to construct relevant alternatives to the timeless-necessity explanation of observed regularities. How, exactly, might we cash out the notion of 'time-limited necessity'?

The first possibility is that, while *N* has related *F* and *G* up to some time *t*, *N* simply stops relating *F* and *G* after *t*. After all, Armstrong hold that it is a contingent matter that *N* happens to relate *F* and *G* (when in fact it is a law that all *F*s are *G*s); there are possible worlds where *F* and *G* are merely accidentally correlated, or not correlated at all, just as it is a contingent matter that Everest stands in the *taller than*

relation to K2. In the latter case, we can obviously imagine that in fact Everest does not *timelessly* stand in that relation to K2 at all: it's entirely possible that suitably cataclysmic shifts in the tectonic plates will eventually render K2 taller than Everest. In other words, it may not even be contingently *true* that Everest is (timelessly) taller than K2; it may only turn out to have been true for some limited period of time. Similarly, I claim, *F* and *G*, while having been related by *N* up to time *t*, might simply stop standing in that relation after *t*.^{viii}

One might reasonably object that the analogy is a bad one: for Everest and K2 to change with respect to the *taller-than* relation, at least one of them has to change its height. But universals do not change. Their instantiations are literally identical: the *F*-ness and *G*-ness that were instantiated last Tuesday are exactly the same as the *F*-ness and *G*-ness that were instantiated in 1543. So how can they bear any relation to one another at one time but not at another?

Well – and this is admittedly rather fanciful, but I'm not sure how else to capture the relevant thought – imagine God watching the Universe unfold. At the beginning of time, he decides it would be nice for all *F*s to be *G*s, at least for the first few billion years, and the easiest way to guarantee that is to glue the universals *F* and *G* together. (God sets a high premium on the Universe being law-governed, but he really can't be bothered fixing the initial conditions and the other laws in such a way that they together deliver the result that all *F*s are *G*s for the first few billion years. That would way too complicated.) So he glues *F* and *G* together with *N*. After a few billion years – in 2010, say – he gets a bit bored with the tedious regularity with which *G*s follow *F*s and he decides to make a change. ('Those humans are getting a bit blasé about this whole science business', he thinks. 'They think they've cracked the secrets of the Universe. Well I'll show them who's boss!') He removes the glue

and – presto! – things down on Earth start getting really unpredictable, causing some serious confusion amongst the scientific and philosophical communities until they eventually get their heads around the fact that the laws of nature have changed. I don't see why God couldn't do this. After all, as I have said, necessity, on Armstrong's view, is contingent: it glues things together that are not glued together in other possible worlds. So why could God not *actually* separate them?

One might be tempted to say at this point that the very notion of a time-limited universal is incoherent. After all, N isn't just any old relation; it's *necessity*, for goodness' sake! It surely wouldn't be a genuine necessary connection if it related F and G at one time but not at another. Well, fair enough: we can reserve the term 'necessary connection' for the relation N , if there is one, such that if F and G are related by N at one time, they are so related at all times. But this raises a second possibility for time-limited necessity, namely that there are 'necessities' that (unlike N itself) are inherently time-limited. Let N_t be the relation such that if it relates F and G , then any F prior to t is guaranteed to be a G , but that guarantee does not extend beyond t . Or, if you prefer, let N_t be the relation such that at all times at which it relates F and G , F s are guaranteed to be G s; and it relates F and G , by definition, up until time t . Either way, future F s may or may not be G s, but if they are, they are only accidentally so. (So on this picture, God starts out by relating F and G with, say, N_{2010} .)

In fact, we don't even need to appeal to time-limited *necessity* to turn the required trick. Armstrong himself allows for the possibility of what he calls 'cosmic epochs': different stretches of time during which different laws hold. He suggests that we introduce the notion of a 'quasi-universal', which is just like a universal except that it involves essential reference to a particular epoch. So an epoch-restricted law

would ‘relate a certain range of quasi-universals (F s in epoch 13, say) to universals, by a necessitation relation’ (1983, 101). What is interesting about cosmic epochs, in the context of the current discussion, is that no time-restriction is placed on N itself. Instead, the time-restriction is built into one of its relata: the quasi-universal F in epoch 13. So the problem of induction, construed in terms of cosmic epochs, is the problem of justifying the claim that N – the same old timeless N – relates genuine universals and not time-limited quasi-universals. (So on this picture, God starts out by relating, say, F_{2010} and G_{2010} by N .)

One can even extract a second way of characterising the relevant time-limited laws that does not involve appealing to any time-limited universals by invoking Armstrong’s distinction between ‘iron’ and ‘oaken’ laws. This distinction is Armstrong’s attempt to deal with a serious problem with his theory of laws, namely that some law statements have the form ‘all F s are G s – except for those F s that are H s’. In other words, H is some factor that, when present, prevents (as it were) N from doing its usual job of guaranteeing the instantiation of G . The problem for Armstrong is that since N relates universals, the law cannot be $N(F \ \& \ \sim H, G)$, since there are no such things as negative universals, and so ‘ $F \ \& \ \sim H$ ’ cannot refer to a genuine relatum of N . Armstrong’s solution is to claim that in this case, $N(F, G)$ is still a law, but an ‘oaken’ one – and so one that does not in fact entail that all F s are G s. (An iron law is thus a law for which there is no such confounding factor H .)

What is important for our current concern is Armstrong’s insistence that ‘the relation of necessitation, N , is the same in the two cases’ (1983, 150) – that is, in both iron and oaken laws. But now we can consider the possibility that one thing that could play the role of H is some period of time, so that the confounding ‘factor’ is, say, its being before midday on October 18, 2010. Of course, this is not obviously a

universal; nonetheless, it is unclear on what grounds (grounds, that is, that would satisfy an inductive sceptic) we could rule out the possibility that *F*s might stop being *G*s for no better reason than that a particular time has passed. (So on this picture, God sets up $N(F, G)$ as an oaken law, with $t=2010$ or *later* as the confounding factor.)

I conclude that there are no grounds for ruling that the notion of time-limited necessity, in any of the senses described above, is incoherent.

Is timeless necessity a better explainer?

The second option for the defender of Armstrong's attempted solution to the problem of induction is to concede that time-limited necessity is coherent, but to argue that the timeless necessity hypothesis (T) is nonetheless a *better* explanation of observed regularity than is (SF), on the grounds that (T) has the advantage of *simplicity* over (SF). Simplicity, after all, is a widely-acknowledged explanatory virtue, and I have granted the rationality of IBE for the sake of the argument.

At first sight this looks like a promising line of objection, for it looks as though all of the ways of cashing out the notion of time-limited necessity proposed above do indeed look less simple than (T), since all of them introduce a temporal parameter. For example, consider the proposal that *F* and *G* are related by N_{now} , or that the present moment in effect works like a confounding factor (on the model of Armstrong's oaken laws). There is a distinct whiff of arbitrariness here: after all, why postulate the existence of N_{now} , rather than any one of the indefinitely many alternative time-limited hypotheses, such as that *F* and *G* are related by $N_{next\ Tuesday}$? To put the point slightly differently, our time-limited candidates introduce an adjustable parameter – the time at which the necessary connection is supposed to break down – whereas no such parameter is present according to (T).^{ix} And, one might

reasonably claim, absence or minimisation of adjustable parameters counts towards simplicity.

The appropriate response to this objection is to point out that (SF) *itself* contains no adjustable parameters: there is no mention of any specific temporal constraint in the formulation of (SF), since (SF) merely asserts that F and G have been necessarily connected *so far*. This explanatory hypothesis simply leaves it open what, exactly, its truthmaker is; and it is only at the level of the possible truthmakers for (SF) that the worry about simplicity emerges. (SF)'s truthmaker could be the existence of any one of indefinitely many time-limited necessary connections, or indeed it could be a timeless necessary connection, since (SF) makes no claim whatsoever about whether or not F and G will continue to be necessarily connected in the future. Moreover, the former, time-limited connections might be of any of the various kinds canvassed above, involving cosmic epochs (so that N itself is not time-limited but its relata are), an oaken law (so that some time t is itself a confounding factor), or whatever. (SF) remains silent on these issues.

Which, then, out of (T) and (SF), is the simpler hypothesis? There is, I think, no sensible way to answer this question one way or the other. The only difference between the two is that (T) makes a positive claim about the future, while (SF) does not. While this of course makes a difference to the relative predictive strengths of (T) and (SF), I can think of no reason why one should additionally think that it makes a difference to simplicity. Thus, as far as simplicity is concerned, the result is a tie. So (T) is not *the* simplest explanation of past regularity, and so IBE does not, at least not on the grounds of simplicity, warrant the inference to (T) rather than (SF).

How might the defender of Armstrong's solution to the problem of induction respond to this move? I can think of two lines of defence. First, one might insist that

(SF) is explanatorily dubious, on the grounds that it (unlike (T)) has many possible truthmakers, and this somehow gives it some deficiency as an explanation that (T) lacks. Second, one might try to argue that predictive strength is an explanatory virtue, and since (T) has predictive power while (SF) has none, (T) is to be preferred over (SF). I shall deal with these objections in turn.

First, then, the worry about truthmakers. There are two slightly different forms this worry might take. First, one might attempt to claim that the very fact that (SF) could be made true by any of various different ontological scenarios renders it explanatorily deficient. Or, second, one might concede that this by itself does not make for explanatory deficiency, but claim that the fact that all but one of those possible truthmakers (the exception being the timeless necessity hypothesis) involve an adjustable parameter makes (T) better than (SF) on the grounds of simplicity after all: the lack of adjustable parameters in (SF) itself has in some sense been achieved merely by suppressing the adjustable parameters that nearly all of (SF)'s truthmakers possess.

The first version of the worry seems to me to require far more from explanations than we in fact require from them. Here's a toy analogy. There are twenty balls in a bag, all of which (unknown to you) are different shades of red. You pull a ball from the bag, and you want to know why you pulled out a red ball. My answer: *all* the balls are red (so whatever you did, you were bound to pull out a red ball). Of course, there are many, many different possible truthmakers for the fact that all the balls are red, since there are many, many precise shades of red, and many, many ways in which those shades might be distributed amongst the balls. This fact about the many possible truthmakers for 'all the balls are red' seems to me in no way to impugn my answer to your question. Indeed, were I instead to have told you

exactly which shade of red each ball was, you might legitimately have complained that I was giving you information that was entirely irrelevant to explaining what I wanted explained. The precise information that ball 1 was maroon and ball 2 was scarlet and ... plays no useful role in explaining why the ball you pulled was *red*. Similarly, I claim, in the case of (SF). If you want to know why all *F*s have been *G*s *so far*, and I tell you it's because *F*s and *G*s have been necessarily connected *so far*, I tell you something that is completely neutral between various different possible facts about what, exactly, makes the claim true. But so what? Maximally specific information about the ontological ground of my claim wasn't what you asked me for, so it is hard to see why giving you such information would enhance my explanation. Of course, it's always nice to know things about the fundamental constitution of reality. But it doesn't follow that precise information about the fundamental constitution of reality always constitutes a better explanation of some fixed fact (such as the fact that all observed *F*s have been *G*s, or the fact that you pulled out a red ball) than does information that leaves various different ontological possibilities open. (SF) does just that, to no detriment to its status as an explanatory hypothesis relative to (T).

The second version of the worry concerned the suppression of adjustable parameters, and the response is similar. Here's another toy analogy. You want to know why Liverpool have failed to score against much weaker teams so far this season. I tell you it's because Torres has been injured and so out of the team. That's an answer that suppresses adjustable parameters in something like the same way that (SF) does, in that my answer leaves it open whether Torres will be back in the team next week, next month, next season, or never. But again, so what? You didn't ask me when the situation was likely to improve. You might have an interest in that question, but providing you with an answer to a question you didn't ask me would in no way

improve on the explanation I gave for the fact that you asked me to explain. Similarly, (SF) leaves it open whether the necessary connection – and so the regularity – will continue for the next week, the next month, until the beginning of the next football season, or forever. But the explanandum includes no claim one way or the other about whether the regularity will continue, so it is entirely appropriate that the explanans doesn't either.

A related concern – though not one that explicitly appeals to simplicity – is raised by Foster in response to the kind of strategy I have been pursuing. He says:

... on this point, I think, the defender of [the 'nomological-explanatory solution' to the problem of induction] can stand his ground. For it seems to me that a law whose scope is restricted to some particular period is more mysterious, inherently more puzzling, than one which is temporally universal. Thus if someone were seriously to propose [a time-limited law] as the correct account [of the regularity], our response would be to ask *why* the relevant law should be time-discriminatory in that way. Why should a certain moment have this unique significance in the structure of the Universe ...? Barring the postulation of a malicious demon, these questions are unanswerable ... we are left feeling that, as hypothesized, nature would be inherently puzzling, and would preclude an explanation of our empirical data which was both correct and, from the standpoint of our rational concerns, fully satisfactory. (Foster 2004, 71)

Foster is here considering the possibility of some *specific* time-limited law being proposed as a rival explanation of the observed regularity; and of course this is not what (SF) does. Nonetheless, one might still worry that the question about a certain moment having a unique significance is still pertinent, given that all but one of the

possible truthmakers for (SF) appeal to the existence of such a moment.

The appropriate response, it seems to me, is to question Foster's 'inherently more puzzling' claim. First, as I just said, (SF) does not positively claim that there *is* one moment that has 'unique significance in the structure of the Universe'. Rather it merely leaves open the *possibility* that there is such a moment (which could be next week, in 4026, ...). For that matter, (SF) leaves open the possibility that there are in fact many such moments; for all (SF) says, it might turn out that Universe starts exhibiting frequent – indeed perhaps even regular – changes in 'cosmic epochs'. Second, we might indeed wonder why the relevant law should be time-discriminatory; but we might equally wonder why a law is *not* time-discriminatory. Of course, psychologically speaking, we all expect the current regularities to persist. A time-discriminatory law thwarts our expectations, and so appears puzzling. But – by the sceptic's lights – we have no epistemic entitlement to the expectations we find ourselves with, and so no entitlement to be more puzzled by any of (SF)'s time-limited truthmakers than by (T). To put the point another way, one might indeed think that (SF) – or one of its possible truthmakers – would preclude an explanation of the data which was fully satisfactory 'from the standpoint of our rational concerns'. But whose rational concerns are these? The primary rational concern of the sceptic is not to make the mistake of thinking that we have any reason to suppose that the future will resemble the past, in the absence of a convincing argument to the contrary. (SF) satisfies *this* concern rather well. It might not satisfy the concerns of those of us who happily set sceptical problems to one side; but from the sceptic's point of view, those concerns are, precisely, *not* 'rational' concerns at all. By 'our rational concerns' Foster seems to mean 'the rational concerns of those of us who don't worry about the problem of induction' – which of course simply begs the question against the

inductive sceptic.

The second line of defence advertised above on behalf of the defender of Armstrong's argument is to appeal to predictive power as an additional explanatory virtue, since, if this is a virtue, clearly it is a virtue that (T) possesses and (SF) lacks. My response is to question whether, in the current context, predictive power should be seen as an explanatory virtue. Armstrong is in the business of trying to provide a sceptic-busting argument for the rationality of induction, and I granted the legitimacy of IBE for the sake of the argument in order to show that, even granted that assumption, Armstrong's argument fails. Of course, granting the legitimacy of IBE involves granting that certain features count as genuine explanatory virtues – simplicity, for example. Now, should predictive power also be granted that status? It can certainly be granted that predictive power counts as an explanatory virtue in the sciences: an explanatory hypotheses that generates novel and interesting predictions is to be preferred to one that merely explains the phenomenon under investigation in an *ad hoc* way that generates no novel predictions. And rightly so: prediction is part of the *point* of science, both in a practical sense and in the theoretical sense that predictive success or failure is the primary arbiter in disputes between theories. But scientific explanation is not our current business: the explanations of the observed phenomena that we are canvassing fall under the scope of metaphysics and not science. And prediction is *not* part of the point of metaphysics, in either a practical or a theoretical sense: we do not, by and large, expect metaphysical theories to generate testable consequences, any more than we expect them to help us build better bridges or cure cancer. So the claim that predictive power is a legitimate explanatory virtue in a metaphysical context is highly controversial, and certainly cannot be inferred to be a virtue merely on the grounds it is a virtue in the sciences.

Moreover, Armstrong's opponent is the inductive sceptic, and, if an argument is to be had about whether predictive power is an explanatory virtue, it is pretty clear which side of the fence the sceptic will be on. And this would be no *ad hoc* manoeuvre, since of course it is precisely the rationality of prediction that the sceptic questions. So to insist that predictive power is an explanatory virtue is to insist upon something that the sceptic takes herself to have good reason to deny. The inductive sceptic holds that, pending a good argument to the contrary, a hypothesis that makes predictions is *eo ipso* a hypothesis that we have no grounds for believing. The argument to the contrary that is being offered – that there are grounds for believing (T) because it is the best explanation of past regularity – turns out to rely on the assumption that predictive power is an explanatory virtue. But this is an assumption which, in the context of IBE, directly entails the denial of the claim that a hypothesis that makes predictions is *eo ipso* a hypothesis that we have no grounds for believing, since it amounts to the claim that an *explanatory* hypothesis that makes predictions *can* be a hypothesis that we have grounds for believing. So in the absence of any argument to the contrary – and here we have an assumption, not an argument – there is no way to persuade the sceptic that explanatory hypotheses are a special case. Hence the necessitarian's argument begs the question against the sceptic.

I conclude that Armstrong's attempt to argue that necessitarianism renders the problem of induction soluble are fatally flawed. In particular, it is either straightforwardly false that (T) provides a better explanation of observed regularity than (SF) does, or else (T) does provide a better explanation, granted an assumption that the sceptic takes herself to have good reason to deny, namely that predictive power is an explanatory virtue – in which case the necessitarian's argument simply begs the question against the sceptic.

5. Ellis's scientific essentialism

Brian Ellis (2001, 2002) calls his overall metaphysical position 'scientific essentialism'. Here's the basic idea. To be a member of a natural kind (an electron, a water molecule, a carbon atom) is to be intrinsically and essentially disposed to behave in certain kinds of ways in certain circumstances. The laws of nature tell us how, in virtue of being the natural kinds of things they are, things are essentially disposed to behave. Electrons are, by their very nature, disposed to repel each other. So it's a law that they do. Moreover, any possible world that contains electrons will be a world in which electrons repel each other: the nature of an electron guarantees that it will behave the same way given the same conditions.

Ellis sometimes characterises his view by saying that, according to him, the laws of nature are metaphysically necessary. But it's important to note that there is an implicit qualification to be made. He doesn't mean (contrary to what the slogan might suggest) that every possible world has exactly the same laws of nature. What he means instead is that any two possible worlds *of the same natural kind* have exactly the same laws of nature. And – roughly – two possible worlds will be *of the same natural kind* if they *contain* the same natural kinds. For example, any possible world with the same kinds of elementary particle as the actual world will have the same chemical elements, compounds and so on; and these things will all behave in the same way as they do in the actual world. In any world that has electrons in it, it will be true that electrons repel each other; in any world where salt and water exist, it will be true that salt dissolves in water, and so on (see Ellis 2001, 249-53).

Ellis explicitly addresses the problem of induction (2001, 283-90; 2002, 134-

7). He says that essentialism:

promises to transform our thinking about scientific rationality and the theory of inductive reasoning. If one believes, as Hume did, that all events are loose and separate, then the problem of induction is probably insoluble. Anything could happen. But if one thinks, as scientific essentialists do, that the laws of nature are immanent in the world, and depend on the essential natures of things, then there are strong constraints on what could possibly happen. (2001, 283)

At first sight, Ellis's essentialist brand of necessitarianism seems immune to the kind of move I offered above in response to Armstrong's necessitarianism. That response depended on the fact that it makes sense to suppose that *F* and *G* might have been necessarily connected in the past, yet fail to be so in the future. But essentialism appears to remove this possibility; after all, if part of what it *is* to be an *F* is to be disposed to produce *G* – if having that disposition is part of the nature of *F*s – then *F*s cannot fail to continue to be followed by *G*s in the future. As Ellis says: 'If there is good reason to believe that something is a member of a natural kind, and good reason to think that it [the natural kind] has such and such a nature, then there is good reason to think that everything of that kind must have this same nature' (2002, 135).

Ellis goes on to note that the essentialist is nonetheless confronted with problems concerning how natural kinds are to be identified and their causal powers revealed. 'But', he says, 'these are the kinds of doubts and concerns that working scientists are accustomed to, and know how to handle. They are not irresolvable sceptical doubts like those generated by Humeanism' (2002, 136).

But Ellis's claim that scientific essentialism removes 'irresolvable sceptical doubts' looks suspiciously like a sleight of hand. He is right to say that for the

essentialist, the epistemological problem concerns our claim to know about the nature and powers of natural kinds; but, I shall argue, this gives us no grounds for thinking that irresolvable sceptical doubts thereby somehow drop off the agenda, so that we are left only with questions that scientists ‘know how to handle’.

Ellis’s basic thought, I take it, runs something like this. Grant that if something walks like a duck, swims like a duck and quacks like a duck, we have good reason to believe that it’s a duck, since its being a duck explains why its behaviour is duck-like. Grant also that *duck* is an Ellisian natural kind, and so has an essence *E*, which not only causally explains why ducks walk, swim and quack in the way that they do, but also explains, and so predicts, a range of other features of duck. (Ducks like eating corn, have webbed feet, breed with other ducks in suitable circumstances, and so on.)^x So, when faced with a suitably duck-like entity, I am warranted in believing that it is a duck – a member of a kind with essence *E* – by IBE: its being a member of the kind with that essence explains why it walks and swims and quacks in the way that it does. And I am then in a position to make a range of other predictions about my duck, since those predictions are licensed by the fact that my duck has essence *E*.

This line of argument shares its basic structure with Armstrong’s: we have a two-step argument, the first of which takes us from observed facts to the existence of some kind of regularity-guaranteeing feature of reality (in Ellis’s case a natural kind with an essential nature rather than Armstrong’s *N*), and the second of which takes us from the existence of that regularity-guaranteeing feature to a prediction that it entails. As with my objection to Armstrong, the question I want to focus on is whether membership of the relevant Ellisian natural kind is, by the sceptic’s lights, the *best* explanation of the observed facts.

To bring out the analogy with Armstrong’s argument, let’s recast the

essentialist argument against inductive scepticism in terms of regularities.^{xi} Thus ('SE' for 'scientific essentialism'):

(SE) All observed *F*s have produced *G*s. The best explanation of this is that the *F*s are members of a natural kind *K*, whose essence is or includes the disposition to produce *G*s. Hence all *F*s (by virtue of membership of kind *K*) produce *G*s.

Let's grant for the sake of the argument, as I did with Armstrong, that membership of an Ellisian natural kind is the best explanation of regularity *simpliciter* – that is, of both and unobserved regularity. The question, again, is whether, when we seek to explain merely observed regularity, another, equally good explanation is in the offing.

Recall that in the discussion of Armstrong, I granted that the appeal to necessary connections does do some explanatory work when explaining observed regularities; it was just that the claim that the necessary connections will *continue* to hold, as entailed by (T), confers no explanatory advantage over the weaker claim that they have held so far (SF). And I argued that the claim that the non-timeless necessary connections that are required for (SF) to be a genuine explanatory alternative to (T) really are conceptually and metaphysically possible. One might think that an Armstrong-style argument can be advanced using the essentialist view of laws, and that this would circumvent the objection, since no equivalent sceptical alternative is available: natural kinds cannot change their essences over time, and so if the existence of natural kinds (along with their dispositional essences) explains observed regularity, no coherent rival explanation involving changes in dispositional essence can be formulated. I shall argue that the objection cannot be circumvented in this way, because relevant sceptic-friendly alternative explanations can in fact be given, the

immutability of natural kinds notwithstanding.

Note that ‘the *F*s’, in our alleged best explanation in (SE), is ambiguous between ‘the observed *F*s’ and ‘all *F*s’. So, disambiguating, we have two candidate explanations for our observed regularity:

(SF*) *The observed F*s are (or were, at the time at which they were observed) members of an Ellisian natural kind *K*, whose essence is or includes the disposition to produce *G*s, and

(T*) *All F*s are members of an Ellisian natural kind *K*, whose essence is or includes the disposition to produce *G*s.

What we need to know is whether (T*) is a better explanation of our observed regularity than is (SF*), and hence whether the inference to (T*) is indeed sanctioned by IBE.

In fact, we need to consider two cases separately, which differ according to whether or not classifying something as an *F* automatically guarantees that is a member of kind *K*. For example, I might want to know why all previously observed salty-tasting, white, crystalline substances found in my salt cellar have dissolved in water, or I might want to know why all previously observed samples of *salt* have dissolved in water. Either way, I shall argue, no solution to the problem of induction is in the offing.

Let’s start with the case where classifying something as an *F* does not automatically guarantee membership of a particular natural kind. We can think of *F*-ness, as characterising the ‘nominal essence’ of the *F*s, in such a way that, in

principle, an object or substance could have *F* but be a member of a different natural kind to the observed *F*s, or perhaps not be a member of any natural kind at all. For example, *F*-ness might here be being white, crystalline, and edible, with a salty taste, and *K* is the kind *salt* (in the restrictive sense of ‘salt’ that just means ‘sodium chloride’).^{xiii} Clearly in principle there could be (and indeed are, in this case) other things that are *F* that are not members of the kind *salt*. Let our explanandum be the fact that all observed samples of *F* have dissolved (*G*) when stirred in water in appropriate circumstances.

Let’s grant that the past observed regularity really is best explained by membership of Ellisian kind *K*. (The fact that all previous *F*s have dissolved in water is best explained by the fact that they were all samples of salt, given that it is part of salt’s essence that it dissolves in water.) The question is, which specific explanation, out of (SF*) and (T*), is the best explanation? To gain any purchase on the problem of induction, the answer has to be (T*), else (SE) fails: that *observed F*s are (or were) members of *K* licenses no inference to the behaviour of *all F*s. But now a familiar question emerges: why should we think that (T*) is a better explanation of the observed regularity than (SF*) is? After all, the only virtue (T*) would seem to have over (SF*) is predictive strength; and, again, I have argued already that predictive strength cannot be assumed to be an explanatory virtue in the context of solving the problem of induction.

The claim that (SF*) is at least as good an explanation of observed regularity as – and does not entail – (T*) need not trade on any curious metaphysical hypotheses analogous to the time-limited universals discussed earlier – metaphysical hypotheses that one might suspect are ruled as incoherent, given the immutability of Ellisian natural kinds. The sceptical possibility we need to entertain is merely that different

natural kinds can share the same nominal essence – something that is not at all bizarre. Think of jadeite and nephrite: their routinely-observable features are so similar that they are both classed as ‘jade’ in ordinary language, but a suitably-equipped chemist can tell them apart: they can empirically verify that they have different underlying chemical compositions.

Of course, such cases are doubtless the exception rather than the rule: when I put the salt-like substance from my salt cellar into water and stir it up, I expect it to dissolve. But the sceptic’s worry, put in essentialist terms, is whether she has any grounds for thinking that the stuff in her salt cellar really is salt, and not some other, perhaps previously unencountered substance with the same observable features, which lacks salt’s dispositional essence. This sceptical scenario is, admittedly, a somewhat far-fetched one – but then, sceptical scenarios usually are. What’s important is that the scenario is entirely consistent with Ellis’s metaphysics.

To make this point a bit more vividly, consider the sceptical possibility that all or most of the regularities that we have observed so far break down at some point in the future. One might be inclined to think that it is precisely this kind of possibility that is ruled out by the existence of Ellisian natural kinds. This would be a mistake, however, because nothing in Ellis’s metaphysics rules out the possibility of wholesale changes in *which* natural kinds there are. Perhaps, for example, the fundamental particles will start behaving in totally different ways. In that case, there will no longer be the *kinds* of fundamental particle that there previously were; there will be new kinds, with new dispositional essences, and the old kinds will have gone out of existence (or at any rate will no longer be instantiated). In other words, the immutability of natural kinds and the corresponding metaphysical necessity of the laws do not guarantee that *which* natural kinds are instantiated remains constant over

time. Something that is a duck today might continue to walk, swim and quack like a duck tomorrow, and yet fail, as of tomorrow, to have the dispositional kind-essence *E* required for duckhood. It would then, of course, no longer be a duck. But nothing in Ellis's metaphysics rules out this kind of change in kind-essence.

One might object that such a possibility *is* ruled out if the relevant essences are essences not just of kinds but of the individuals that comprise the kind, so that the very same particle (say) cannot lose the disposition in virtue of which it is a member of a given kind, because to lose that disposition would destroy the very identity of the particle. Ellis himself is 'reluctant to accept that the individual essence of a thing belonging to a natural kind includes its kind essence' (2001, 238). Still, one could certainly hold the view that Ellis is reluctant to accept. Unfortunately, that would not help. As Ellis says, on this view, 'an individual of one kind could not possibly be transformed into something of another kind, although it might cease to exist and be replaced by something else' (*ibid.*). So we can simply recast the sceptical possibility as a kind of wholesale replacement of one set of entities by another.

The idea of 'replacement' might seem – and indeed is – metaphysically curious. But it is not quite as curious as it might seem at first sight. Suppose that Bert is essentially a person, and that certain psychological features are essential to being a person. Then there are two ways in which you might replace Bert (a person) with a corpse (not a person). One way would be to forcibly remove him from the room, and put in his place a dead body. The other would be simply to kill him: this would ensure that Bert goes out of existence, and a corpse comes into existence, and hence that Bert has been 'replaced' by a corpse. It is this second sense of 'replacement' that is at work in the suggestion that one set of entities – entities that are essentially members of a given natural kind, and hence are essentially disposed to behave in certain ways –

might be replaced by another set of entities: entities that are essentially members of a different natural kind, and hence are essentially disposed to behave in different ways.

To sum up: that the *so far* observed *F*s have produced *G*s is explained just as well by (SF*) as it is by (T*). And (SF*) is compatible with future *F*s failing to produce *G*s, since it is consistent with the hypothesis that future *F*s will be members of different natural kinds (or indeed members of no natural kind at all), such that neither their individual essences, nor their essences *qua* members of any natural kinds they are members of, do not include the disposition to produce *G*s. Hence the inference to ‘all *F*s produce *G*s’, via IBE, is blocked.

Let’s turn now to the other case, where classifying something as an *F* *does* automatically guarantee membership of a particular natural kind. For example, imagine that I want to know why all the electrons I have observed have been repelled by positively-charged particles. Here, the natural kind just *is* the kind, *electron*. (This need not impugn the explanatory status of the claim that the electrons I’ve observed are all members of a natural kind (*viz*, the kind *electron*), since this might not be something I already knew.) In that case, the relevant explanatory hypothesis can legitimately be claimed to be (T*), as required, rather than (SF*): given that any electron, by definition, will be a member of that natural kind, the question of whether unobserved as well as observed electrons are members of the kind and so have the relevant disposition is guaranteed to get the answer ‘yes’.

Does this provide a solution to the problem of inductive scepticism?

Unfortunately not. Note, for starters, that it is a part of the essentialist thesis that claims of the form ‘all *F*s produce *G*s’ are metaphysically necessary, if the disposition to produce *G*s is part of the essence of *F*s. So no appeal to IBE is required in order to establish the truth of the general claim. Indeed, it looks as though what is being

explained, in this case, is not so much why all previously observed *F*s have produced *G*s, but rather why it was right to classify them as *F*s in the first place. In essence, the case is no different to being asked to explain why all previously-encountered samples of water were composed of H₂O molecules. All that can be said in response, on an essentialist view, is that being so composed is just what it *is* to be water: if the previously-encountered samples had not been so composed, they would not have been samples of water.

To put it another way, either the person requesting the explanation knows what the essence of *F*s is, or she does not. If she already knows, then there is nothing that needs to be explained: she already knows that all *F*s produce *G*s, so the question of why the *observed F*s have produced *G*s does not need to be asked. If she does not know the essence of *F*s, then the question is legitimate; but the answer merely tells her what that essence is.

Still, to know what the essence is, is to know that all *F*s produce *G*s. Isn't that what we wanted in order to defeat the sceptic? Well, no, because – again – to know this is not, by the sceptic's lights, to know anything about the future. As we have seen, it is an entirely coherent metaphysical possibility, even given Ellis's metaphysics of natural kinds, that there is a radical change in which natural kinds there are. So knowledge that all *F*s produce *G*s will only license inference to some future fact (in particular the fact that a *G* will be produced) on the assumption that there will be some more *F*s. And this is something that cannot be established by inference to the best explanation, when the explanandum in question concerns only what has happened in the past. *If* I know that a given subatomic particle is an electron, then I know how it is disposed to behave. But – the sceptic will ask – on what grounds do I assert that the particle is indeed an electron, as opposed to, say, a *schmelectron* – a

hitherto unencountered particle that behaves in a completely different way?

6. Conclusion

In essence, Ellis's proposed solution fails for the same basic reason as Armstrong's. If we are trying to solve the problem of induction by appealing to IBE, it needs to be the case that the IBE being invoked does not itself require us to take some illicit inductive step. But in both Armstrong's and Ellis's cases, there *is* an illicit inductive step, in that each presupposes that there is no available alternative explanation that does not go beyond what has been observed (or beyond the so-far-underlying nature of what has been observed). Armstrong assumes that the only available kind of necessity is the timeless variety. Ellis assumes, in effect, not only that natural kinds have immutable natures, but also that *which* natural kinds there are is also immutable. So the attempt to run an Armstrongian argument on the basis of Ellis's metaphysics runs up against a similar problem to that faced by Armstrong's own argument: it fails to consider the possibility that there might be a change in which natural kinds there are.

One might object to the line of thought I have been pursuing by claiming that I have been misinterpreting Armstrong's and Ellis's position on the problem of induction. It could be argued that they are not really intending to *solve* the problem of induction; rather, they are making the weaker claim that the necessitarian has the resources for legitimating inductive inference while the Humean does not. In other words, *if* one accepts a worldview according to which there are timeless necessary connections or dispositional essences of eternally-existing natural kinds, then one has grounds for believing that the future will resemble the past; whereas if one believes in no such things, one lacks the grounds for inductive inference. Thus Stephen

Mumford, in his exposition (and apparent endorsement) of Armstrong's argument, sketches it as follows:

[R]egularity theories are left with the problem of induction. Because they grant no inner connection between being *F* and being *G*, they have no basis on which to infer from observed cases to unobserved cases. All known things that are *F* may be *G*, but that does not support the inference that all things that are *F*, whether observed or not, are *G*. In contrast, a realist about laws might claim that there is an inner connection, which provides a reason to think that unobserved cases will be like the observed ones. (2007: 45)

But if this is what Armstrong and Ellis intend, then they are invoking an illegitimate double standard. Of course, realists themselves do not face the problem of induction, in the sense that they believe in something (namely timeless necessity) that delivers a guarantee that chaos will not descend. But Humeans typically believe in something that does the same job, namely the timeless regularity of nature. (Our 'realist' here is someone who believes not only that the realist analysis of lawhood is the right analysis, but also that there are laws, so analysed. Similarly, our Humean here is someone who believes not only that a regularity-based analysis of lawhood is the right analysis, but also that there are laws, so analysed.) The Humean's belief provides an excellent reason to think that past regularities will persist into the future – indeed, as good a reason as the realist has. Of course, the realist might retort that the Humean has no grounds for this belief. But that takes us back to the original argument. If we are arguing about the grounds for belief in the existence of timeless laws (or immutable natural kinds), then the Humean is entitled to respond along the lines I have been suggesting: the Humean and the realist alike appear to lack non-inductive grounds for belief in such laws (kinds), so they are in the same boat.

At this point, the realist *would* be entitled, at least *prima facie*, to point out that the realist's position is better, in that only the realist can offer an *explanation* for the timeless regularity of nature. But whatever the merits of this response, it has nothing whatsoever to do with the grounds for inductive inference. Perhaps, once the timeless regularity of nature is agreed on all sides, the Humean is guilty of failing to explain it. But she is no more guilty than the realist of susceptibility to the problem of induction.

Bibliography

- Armstrong, D. M. (1983) *What is a Law of Nature?* Cambridge: CUP.
- Beebe, H. (2006) 'Does Anything Hold the Universe Together?', *Synthese* 149, 509-33.
- Blackburn, S. (1993) 'Hume and Thick Connexions', in his *Essays in Quasi-Realism*, Oxford: OUP, 94-107.
- Dretske, F. (1977) 'Laws of Nature', *Philosophy of Science* 44, 248-68.
- Ellis, B. (2001) *Scientific Essentialism*, Cambridge: Cambridge University Press.
- Ellis, B. (2002) *The Philosophy of Nature*, Chesham: Acumen.
- Foster, J. (1982-3) 'Induction, Explanation, and Natural Necessity', *Proceedings of the Aristotelian Society* 101, 145-61.
- Foster, J. (2004) *The Divine Lawmaker*. Oxford: Oxford University Press.
- Mellor, D. (1995) *The Facts of Causation*, London: Routledge.
- Mumford, S. (2007) *Laws in Nature*. London: Routledge.
- Strawson, G. (1987) 'Realism and Causation', *The Philosophical Quarterly* 37, 253-77.
- Strawson, G. (1989) *The Secret Connexion*, Oxford: Oxford University Press.
- Tooley, M. (1977) 'The Nature of Laws', *Canadian Journal of Philosophy* 7, 667-98.

van Fraassen, B. C. (1980) *The Scientific Image*, Oxford: Clarendon Press.

van Fraassen, B. C. (1989) *Laws and Symmetry*, Oxford: Clarendon Press.

Endnotes

ⁱ Thanks to the audiences and individuals who have heard and comments on previous versions of this paper, including Alexander Bird, Chris Daly, Brian Ellis, Philip Goff and Barry Loewer. Thanks also to a referee for this journal for some very helpful suggestions, and to both the Leverhulme Trust and the AHRC, whose support – in the form of a research fellowship and the funding of the ‘Metaphysics of Science’ project respectively – made the completion of this paper possible.

ⁱⁱ John Foster also argues that realism about laws solves the problem of induction (and indeed his original argument came out at around the same time as Armstrong’s; see Foster 1982-3). However Foster argues that the two brands of necessitarianism just listed fail to provide viable accounts of the nature of laws and proposes instead a theistic analysis of laws, according to which the laws express the intentions of ‘the divine lawgiver’; see his 2004. So Foster might be characterized as endorsing a rather unusual brand of necessitarianism.

ⁱⁱⁱ See for example Strawson 1987 and 1989 (Chapter 5). For Strawson, it is Causation (with a capital ‘C’), rather than the laws of nature, that explains the regularity of the Universe.

^{iv} See Mellor 1995, 75-6.

^v I pursue this line of thought in my 2006.

^{vi} For van Fraassen's scepticism about IBE in the context of inference to the unobservable, see his 1980, Chapter 2. For his argument against the rationality of IBE in the context of explaining regularities, see his 1989, Chapter 6.

^{vii} Actually, Armstrong's own view of laws doesn't have as a consequence that ' $N(F,G)$ ' entails 'all F s are G s', because he thinks that the entailment doesn't go through for 'oaken' laws (see Armstrong 1983, 147-50). I'll let that pass.

^{viii} Note that this possibility is not ruled out just by Armstrong's claim that universals in general – and so N in particular – *exist* timelessly if they are instantiated at all. That claim is secured, Armstrong thinks, by the mere instantiation, at any time, of a given universal: a universal that is instantiated at t but not at any time thereafter does not go out of existence; it simply fails to be instantiated after t .

^{ix} Thanks to an anonymous referee for making this point, and for making me rethink the whole argument of this section.

^x It might be argued that *duck* isn't a very good example, both because Ellis thinks that biological species are not natural kinds and because, even if they are, it is controversial what their essences are. These details are irrelevant, however; the same general point could be made with chemical elements, subatomic particles, or whatever.

^{xi} In the duck example just given, we are explaining merely why *this* duck-like entity manifests its duck-like features, rather than why *all* previously observed duck-like entities do. Correspondingly, the aim in the example was to infer things about the future behaviour of this particular entity, and not of duck-like entities in general. However, nothing hangs on this difference, so I shall stick to the observed-regularity-to-universal-regularity version of the argument.

^{xii} Of course, this is not sodium chloride's nominal essence in Locke's original sense, since for Locke, nominal essence by definition characterizes a kind.