

Regularity Theory and Inductive Scepticism The Fight Against Armstrong

Benjamin Smart

Introduction

David Armstrong argues that given the Humean sees “the true form of an inductive inference (to be) simply an inference from the observed cases to the unobserved cases. And, given that the law is just the observed plus the unobserved cases, *that* inference,... is an irrational inference” (1983: 53)¹. The Humean, according to Armstrong, is committed to inductive scepticism².

I show in this paper that in fact the opposite is true. The regularity theorist is entirely committed to non-scepticism about induction, and furthermore that contra Armstrong, the regularity theorist is fully entitled to this position regardless of whether he’s justified in inferring the universally quantified conditionals he posits as laws. To achieve this aim I postulate a ‘regularity-relation between universals’; a relation that plays the same role as Armstrong’s Natural Necessitation Relation, but entails a high probability of our conclusions about the unobserved being correct (less than 1 but certainly sufficient to justify inductive inferences) purely through observed regularities and a priori mathematics, without needing to bring any extra, ‘spooky’ entities into our ontology.

¹ *Ibid.*

² Armstrong actually attributes this position to Hume himself, but whether or not Hume genuinely believed induction to be irrational is debated. Regardless of the outcome of this debate, this paper is concerned with what the ‘Humean’, or perhaps rather less ambiguously the ‘regularity theorist’ should believe, as opposed to what Hume himself believed.

Regularity Theory and Inductive Scepticism

1. Hume and Induction

Although this paper is primarily concerned with *laws of nature* and not specifically *causal laws*³, the best indication of the regularity theorist's commitment to inductive non-scepticism is found in Hume's discussion of causation.

When Hume discusses the formation of ideas from impressions, or our identification of causes and effects, he certainly recognises the inductive inferences involved. If Hume were to be (as traditional interpretations imply) a strict inductive sceptic, it seems he should conclude that any attempts to identify causes and effects are futile; this would be a strange opinion to attribute to Hume. After all, Hume spends considerable time outlining the conditions under which causes and effects *should* be identified. As Helen Beebe suggests, "Hume's rules appear to tell us that we *should* seek out hidden causes; but if he is an inductive, and hence causal, sceptic the rules lack any normative force: no purpose is served by acquiring more, or more refined, causal beliefs" (2006:43). It seems to me that if no inductive inferences were in any way justified, no such set of rules would be better than any other.

Hume does go on to say that we can conceive of "a change in the course of nature; which sufficiently proves, that such a change is not absolutely impossible", but this is more an indication of Hume's recognition of the problem of induction, and the lack of *certainty* that ensues from inductive inferences, than it is evidence of his supposed inductive scepticism. It would be perfectly consistent for Hume to make this assertion, and still hold that inductive reasoning is often justified. He could still believe that we can have good reason

³ It is possible to dismiss the regularity theory of causation, whilst accepting the regularity theory of laws of nature, by denying that one can reduce cause to law; that is, by claiming that there is more to causation than nomic connection³. A regularity theorist of laws of nature could accept $\forall x(Fx \rightarrow Gx)$ to adequately represent laws concerning ravens and blackness and so forth, but deny that the same can be done for causation. The regularity theory of causation thus entails the regularity theory of laws of nature, but the regularity theory of laws of nature does not entail the regularity theory of causation.

LYCEUM

to identify a cause-effect relation purely from experience, even if these reasons are non-deductive and fallible.

Whether or not Hume himself was a non-sceptic about induction is inconsequential for the purposes of this paper. What *is* important is that he, as a regularity theorist, *should* have been. I will first show that if a regularity theorist wishes to identify any laws of nature (even with an admission of fallibility), he *must* be a non-sceptic about induction.

2. Why a Regularity Theorist Must Reason Inductively to Identify Laws of Nature

A regularity theorist believes that when we identify a law of nature, we do so by observing a constant conjunction between certain properties (the property of being a raven, and the property of being black, for example). The observed instances are, of course, constrained to our present and past experiences. According to the regularity theorist, the constant conjunctions that make up a law of nature must hold omnitemporally and omnispatially; “there must be a constant union betwixt the cause and effect. ’Tis chiefly this quality that constitutes the relation” (1985: 223). It follows that if the regularity theorist is to justify his identification of laws, he must justify his belief that the constant conjunctions identified will hold across all spatio-temporal regions. He is making conclusions about the entirely unobserved future, from the partially observed past, in so doing committing himself to the rationality of inductive reasoning.

It may be argued, of course, that the regularity theorist need not *identify* laws of nature in order to maintain his primary beliefs about what a law is constituted by. For a universal regularity to hold; that is, for a law to exist, nobody needs to actually know, or even believe that it’s a law at all. Even when a law *is* identified, the regularity theorist must accept he may be wrong. The regularity he thought was a law can always turn out not to be, as the regularity can always break down at some point in the future (or may even have already

Regularity Theory and Inductive Scepticism

broken down at some unobserved point in the past). So whether or not a regularity theorist identifies a regularity as a law has no bearing on whether it actually *is* a law. Nevertheless, it's clear that the regularity theorist attempts to identify laws, so the problem remains.

3. Armstrong's Argument

Armstrong argues⁴ that the rationality of induction is a necessary truth, not just analytically, but for some 'deeper reason'⁵. He proposes that the necessitarian can rationally predict the continuing uniformity of nature by inferring a 'natural necessitation relation between universals', $N(F,G)$, as the best explanation for our observations. According to Armstrong though, there is no way for the regularity theorist to justify inductively derived predictions about the future. He claims their predictions about unobserved events are not grounded by inference to the best explanation, but based solely on the 'pattern of inference: observed instances to unobserved instances'; where e (the observed instances) is inductive evidence for h (claims about unobserved instances), Armstrong suggests the regularity theorist reasons as follows: $e \rightarrow e+h$, $e+h \rightarrow h$, which of course is reducible to $e \rightarrow h$. This is to be regarded as an irrational inference.

In the pattern of inference which I favour (the explanatory law) we have first a passage from observations to the entity which best explains the observations. It seems reasonable to regard this as a rational, although non-deductive, inference. Second, we have a deductive passage from the entity to the unobserved cases. But what makes the Regularity theorist's preferred pattern of inference rational? On his view the law does not explain the observations⁶.

⁴ See Armstrong; 1983: 54.

⁵ The details of which I won't go into, but for further discussion see Armstrong, *What is a Law of Nature* chapter 4.

⁶ I will argue that the law does in fact explain the observations.

LYCEUM

As Hume pointed out, the observed cases do not entail that the unobserved cases will resemble them. There seems to be no other way to explicate the rationality of the inference. (1983: 56)

According to Armstrong then, a necessitarian can rationally infer laws (that are in some way distinct from the observations) that support counterfactuals and “entail conditional predictions about the future”⁷, but the regularity theorist cannot. The regularity theorist must, therefore, be an inductive sceptic.

Armstrong concludes in favour of N(F,G) via inference to the best explanation, but what is this an explanation of? Consider the relation between the universals ‘ravenhood - F’ and ‘blackness - G’. Armstrong’s can’t be an explanation of why *all* ravens are black (omnitemporally), as this ‘fact’ is not available to him prior to N(F,G) coming into the picture. That *all* ravens are black is not what Armstrong is explaining, but an implication of his explanation. What Armstrong is actually trying to explain is why all the *observed* ravens have been black.

It seems then, that Armstrong feels that for the regularity theorist to justify inferring his laws, the laws must “explain the observations” (1983: 56). According to him though, universal regularities, being just the sum of observed and unobserved instances, provide no substantial explanation for the observations themselves. However, by “explain the observations”, Armstrong must have in mind an explanation for *why* the objects have the properties they do. This seems to be asking too much of the regularity theorist, who by the very nature of his theory rejects that any such explanation exists. It seems to me that all the regularity theorist has to provide is an explanation for why the ravens *have been observed* to be black, and this is a very different prospect. If asked, “why have all the ravens I’ve observed been black?”, “because all ravens are black” looks like a reasonable explanation. I will argue that if he can reasonably take the step from “this raven has been observed to be black” to “this raven is

⁷ *Ibid.*

Regularity Theory and Inductive Scepticism

black”, the regularity theorist can justifiably infer that the next raven he observes will be black.

I demonstrated why the regularity theorist must be an inductive non-sceptic in order to derive laws (in the form of universally quantified conditionals), but I will further argue that the regularity theorist need not make *universal* generalizations in order to make inductive inferences, but may use temporally or spatially restricted ‘laws’ in order to make his inferences. As a matter of fact, all the regularity theorist needs to justify his conclusion that the next F will be a G, is that within his population of Fs, all or *most* Fs are Gs.

4. The Law of Large Numbers: Support for the Regularity Theorist’s Right to Reason Inductively

The Law of Large Numbers shows that for any finite population, proportions in large samples are highly likely to resemble proportions in the total population from which the sample is taken, and *ipso facto*, a population is likely to resemble a large sample of that population. If, for example, we choose 3000 random ravens from a population of a million ravens, half of which are black, half white⁸, the probability of the proportions of that sample being within 3% of the proportions of the total population (between 47% black and 53% black) is greater than 0.9, so when we don’t know the proportions in the total population, it is rational to assign equal proportions to the total population as those we get in our sample. As D.C. Stove puts it:

Whatever the proportion of black ravens may be in a population of a million, at least nine out of ten 3000-fold samples of that population do not diverge from that proportion by more than 3% in the proportion of black ravens they contain (1986: 70).

⁸ Note that a 50/50 proportion will give the lowest probability of a representative sample. A population of 100% ravens will of course give a 100% chance of a representative sample.

LYCEUM

The calculation of these probabilities is very simple *a priori* matter. One merely calculates the number of possible 3000-fold samples in the population, and then the number of these possible samples whose proportions of black-to-white ravens fall within 3% of the proportions of the total population (I'll call these samples 'representative samples'). We then divide the latter by the former to get our probability. It is simply a mathematical truth that more large samples are representative samples than non-representative samples⁹, and so for any random large sample of a finite population, one is more likely to get proportions closely resembling the proportions of the total population than not.

It is important to note that once we have a sample greater than 3000, the size of the total population does not have a significant affect on the proportion of 3000-fold samples representative of the population; that is, if we have a total population of a hundred trillion instead of a million, the majority of samples will still be representative. This may sound counter-intuitive, but it is easily shown to be true mathematically. If the law of large numbers argument is sound, it seems that the regularity theorist can justifiably make inferences from the observed to the unobserved, despite Armstrong's claim to the contrary.

Objections have been raised to this attempt at justifying induction on the grounds that we may never be able to observe a truly random sample, and even if we did, it is unlikely that we could know it to be random. Following Campbell and Franklin I don't see this as problematic, but I will address this objection in detail in the next section.

Another objection regularly raised is that many of our inductive inferences are made with respect to infinite populations, as the mathematics used in Stove's argument could not be put into practice (there would be an infinite number of representative and non-representative samples). It seems to me, though, that when making inferences about particular events or instances, we rarely have to consider infinite populations. Let's assume, for the sake of argument, that there are an infinite number of black ravens. To make an inference to the colour of the *next* raven, I only need to consider the population to be all the past ravens plus

⁹ Assuming, for the time being, an equal probability of choosing any possible sample.

Regularity Theory and Inductive Scepticism

the next raven to be observed. The population is no longer infinite, and so the law of large numbers can be implemented. Granted, this does not allow us to make inferences about the colour of *all* ravens, but this is not needed to predict the colour of the next raven. So long as the next raven forms part of the population designated, I am justified in making my prediction. So the regularity theorist does not need to postulate *universal* laws, he can pick out temporally or spatially restricted populations, of which the unobserved raven is a member, and where there is a large enough population and we have large enough sample to apply the law of large numbers.

5. The Randomness Objection

The Law of Large Numbers provides excellent ammunition for the Regularity Theorist against Armstrong, but the mathematics looks only to be applicable under the assumption that each observable has an equal probability of being observed. It seems that this is not the case with the observation of ravens, however. Surely I have a higher chance of observing the raven in my garden, than I have of observing a raven a thousand miles away. What we apparently need is a ‘fair sampling procedure’, and it’s not clear this is available.

Scott Campbell and James Franklin, however, have argued that there is really no need for randomness whatsoever. “In fact, the demand for randomness... leads to nothing but absurdity.” (2004: 83) In short, Campbell and Franklin argue that even if you have good reason to think your sample is not random, you are often still justified in believing your sample is representative of the total population.

Should we, in knowing that all particles observed over the last few hundred years with negative charge have repelled one another, conclude the electrons prior to the first observations attracted one another? Obviously not. Campbell and Franklin conclude that unless we have a genuine reason for believing our sample is not representative “then it is rational to think that it probably is, *because most samples are*. [To believe otherwise] is equivalent to

LYCEUM

holding that we cannot suppose that we are likely to draw a red ball out of barrel of a hundred balls, 99 of which are red, unless we know the method of drawing out the ball is random or unbiased”(2004: 84). Essentially, in our case, all we need to know to justify applying the Law of Large Numbers is that we have no reason for thinking our sample of ravens is not representative¹⁰.

Campbell and Franklin recognise Indurkha’s objection that a sample of ravens in England may not be representative of a sample in some mountainous region (ravens in a mountainous region may be white, for example, for camouflage purposes). However, they respond by showing there is no *a priori* justification for making these conclusions, and so the objection cannot affect the sampling thesis in general. Admittedly, given empirical evidence, we may well have good reason to be cautious in making generalisations about animals outside of our restricted spatial region, but for many of the inductive inferences we make there is no evidence of this kind, and all the regularity theorist needs to show is that he is justified in reasoning inductively *in some instances*. Indurkha’s objection, therefore, makes little head-way in supporting Armstrong’s conclusions.

6. Armstrong’s response to the law of large numbers approach

Armstrong does discuss the regularity theorists’ appeal to probability and the law of large numbers, but dismisses the validity of this approach by introducing the problems posed by ‘unnatural’ predicates. It seems at least possible that our observation of ravens may lead us to believe these ravens are black, when in fact they are ‘bleen’ (black before the year 3000 and green thereafter). In fact, if this objection held, the Law of Large Numbers should force us to make an infinite number of inconsistent conclusions!

The grue problem is one that applies to the rationality of induction in general, but it has been addressed. Goodman¹¹ has provided arguments to

¹⁰ The fact that we may have these reasons is irrelevant here.

¹¹ See Goodman; 1979 ch.III.

Regularity Theory and Inductive Scepticism

suggest we should only accept natural predicates like green and black, purely because these are the best ‘entrenched’ predicates; that is, predicates like green and black are routinely used by the general populous, and ‘unnatural’ predicates like bleen are not. Restricting the regularity theorist’s inferences to these kinds of predicates would solve the problem. However, Armstrong argues that it is “impossible to see how the new principle [of restricting inferences to natural predicates] is to be justified” (1983:58). Although he says nothing more as to why it couldn’t be justified, I assume the reasoning goes something like this: Whereas Armstrong may rule out grue-like predicates, as only natural predicates are, in Goodman’s words, “well-behaved predicates admissible in lawlike hypotheses” (1979:79), the regularity theorist does not hold a lawlike hypothesis, so this kind of response is unavailable to him. However, I believe the regularity theorist too can make the claim “*nothing is grue*¹²!”

It seems to me that the regularity theorist does, in fact, have an adequate response to Armstrong. He can simply appeal to the very principle this objection is supposed to rule out; that is, he can appeal to inductive evidence. There have, as yet, been no *confirmations* of any objects instantiating unnatural colour predicates like grue. As the date of colour change is completely arbitrary, it should be equally probable that the important dates in the unnatural colour predicates’ calendar are in the past, but no such predicates have ever been seen to hold. Armstrong, when articulating his concerns about the grue problem, claims that emeralds may well be grue instead of green, where the change occurs in the year 2000AD. We are now in 2008, and no colour changes were observed in the year 2000AD, so we have confirmation that emeralds are not grue (of course that doesn’t rule out other unnatural predicates where the changes would occur at some future date). The claim is that whenever we have postulated specific unnatural colour predicates in the past, and the dates on which the colours were meant to change has now passed, we have confirmation that these objects were not, in fact, the postulated unnatural colour predicate.

¹² Or any other unnatural predicate for that matter.

LYCEUM

This looks to be good inductive evidence to suppose our future postulations of certain objects instantiating unnatural predicates will also turn out to be false.

An objection is bound to be raised here. I'm using the very principle that the grue problem is supposed to undermine, to answer the grue problem. It is still the case that unnatural predicates may become apparent in years to come. This may seem problematic, but I see the burden of proof to be very much in the hands of those who think the regularity theorist must be an inductive sceptic. I have shown that the regularity theorist may have good justification to reason inductively independently of the grue problem, and given that induction is necessarily rational, if reasoning inductively in conjunction with my justification of induction can help avoid the grue problem, it's not obvious why it shouldn't be allowed.

7. Can a Universally Quantified Conditional be a Good Explanations for why "All Fs Have Been Observed to be Gs"?

Ultimately, my claim is that *the best explanation* for the observed instances being in the proportions they are, is simply that the same (or at least very similar) proportions would be found in the entire population.

Consider an opaque pot filled with a billion marbles. We randomly pick out five million marbles, and they all happen to be black. If asked why all the marbles have been observed to be black, when there are millions of potential colours, the explanation "because *all* the marbles in the pot are black" seems perfectly reasonable. We have observed a large sample of balls from the pot, and as a result can justifiably apply The Law of Large Numbers to make conclusions about the total population¹³, which in turn explains our observations. Even if you deny that we can justifiably assert a universally quantified conditional on the basis that all we have shown is that it's probable that most ravens are black, as opposed to all ravens are black, this conclusion still enables us to make justified predictions about the colour of future ravens.

¹³ This requires an additional step: see below for details.

Regularity Theory and Inductive Scepticism

Nevertheless, in the same vein as Campbell and Franklin, I maintain that we have no positive reason to think that our sample includes non-black ravens on the basis of our observations, so it seems reasonable to suppose, given the evidence, that it doesn't. This conclusion is of course fallible, but such is the nature of inductive inference. Regardless of this issue (the resolution of which has no bearing on whether the regularity theorist can make inductive inferences), there may still be worries about circularity here, and it is this circularity that brings into question whether these conditionals can be good explanations in this context.

When I explain my observing a ball's blackness by appeal to all (or most of) the balls in the pot being black, it looks as though I'm simply stating "this ball has been observed to be black because this ball is black". Although this in itself is not a circular explanation, to get to the proposition "all (or at least most of) the balls are black" from the observation of black balls, the regularity theorist does have to take the step from "this ball looks black", to "this ball *is* black". As already demonstrated, if this move is allowed, we can then infer the blackness of all the balls in the pot (by applying the same principle to every ball we observe, and then using the Law of Large Numbers to make inferences about the total population), and explain our observations in virtue of these inferences.

This is in itself potentially troubling, as it rests on the assumption that we are justified in our move from "this ball looks black" to "this ball *is* black", but one would expect a G.E. Moore-style response should suffice: I observe a ball that looks black; I am more sure that balls that look black are black, than I am of any other possibility; therefore I'm justified in asserting that this ball is black!¹⁴

Nevertheless, "because all the marbles in the pot are black" is still a "self-evidencing explanation" (Hempel 1965: 370-4). As Lipton suggests, however, the circularity of an explanation does not mean it shouldn't be considered an explanation. Indeed, many of the scientific explanations we accept are of

¹⁴ As a matter of fact, Armstrong himself writes that G.E. Moore is justified in vindicating commonsense (1983: 53). It certainly seems commonsense to allow the step from "this ball appears black" to "this ball is black".

LYCEUM

precisely this kind. I might explain my thinking a certain star is moving away from me by the observable ‘red shift’, but my explanation for my observation of red shift would be that a star is moving away from me¹⁵, “...what is significant is that the circularity is benign” (2004: 24). It does not affect the justification of my explanation of why I observed red shift, nor the explanation of why I think the star is moving away from me. I claim that universally quantified conditionals can be benignly circular explanations. The observation of a large sample of black balls allows us to infer that all (or at least most of) the balls in the pot are black, which in turn explains why all the balls have been observed to be black, and allows us to make claims about the colour of balls taken out of the pot in the future. This example is almost identical to the ‘red-shift’ example above, and yet nobody questions the explanatory value there. The circularity does not affect the justification of why I think all the balls in the pot are black, neither my justification of why I think the next ball I observe will be black, so why shouldn’t “because all the balls in the pot are black” be a good explanation?

8. Natural Necessitation Relations and the New Regularity Relation

Armstrong claims that the natural necessitation relation between universals can be found via inference to the best explanation, and that given this ‘law’ is in some way distinct from unobserved instances, inferences about the nature of these unobserved instances can justifiably be made; that is, $N(F,G)$ somehow entails that all future Fs will be Gs. However, I claim that the regularity theorist can use the same methodology to arrive at a relation between universals: $R(F,G)$, where R is the ‘regularity relation’, which holds between universals in our set population (whether this be a universal, or spatio-temporally restricted population).

The regularity relation is easy to comprehend. It is a contingent relation inferred through the regular observation of Fs being Gs, where the Law of Large

¹⁵ *Ibid.*

Regularity Theory and Inductive Scepticism

Numbers can be applied as justification for making conclusions about the total population of Fs. It is nothing over and above the instances themselves, and *a priori* mathematics. $R(F,G)$ holds when these conclusions can *justifiably* be inferred by the law of large numbers. Once $R(F,G)$ has been inferred, in exactly the same way as $N(F,G)$ can stand between observation and conclusions about the unobserved for Armstrong, $R(F,G)$ can stand between observation and conclusions about the unobserved in the pattern of inference for the regularity theorist. We are left with the pattern of inference $e \rightarrow R(F,G) \rightarrow h$, a rational inferences that completely defuses Armstrong's objection.

One may object that the regularity relation need not identify universal regularities, as it can feasibly be applied to spatio-temporally restricted populations. I can allow this, though. The only concession that has to be made is that if we only admit omni-temporal and omni-spatial regularities to be laws, these spatio-temporally restricted regularities would not be laws. That's not to say we're unjustified in making inductive inferences about unobserved instances within that population.

Secondly, one may object that the number of instances needed to justifiably infer a conclusion via the law of large numbers is vague. We do not know the size of our population in most instances, so we can never precisely calculate the probability of our sample being a near population matcher. Furthermore, what probability actually justifies our inferences? Is 80% enough? Do we need 90%? Does a sample of 3000 completely justify our inferences, but a sample of 2999 provides no justification at all? This seems more problematic, but I the objection can be avoided by accepting the notion of degrees of justification. I would be more justified in my belief that 50% of a particular (fair) coin's tosses will land heads after a sample of 10,000 coin tosses (50% of which landed heads) than after a sample of 5,000 coin tosses (50% of which landed heads). I have already shown, however, that with samples over 3000 we have over a 90% chance of our sample being representative (a near population matcher). Given that the regularities we are generally concerned about in nature generally involve a sample of far more than 3000, and that a 90% chance of our

LYCEUM

sample being representative seems more than sufficient, I don't think this vagueness objection should overly concern us. If the objector is still unsatisfied, he should bear in mind that if this is problematic for the philosopher who wishes to infer a regularity relation between universals, it is equally problematic for he who wishes to infer a natural necessitation relation.¹⁶

The natural necessitation relation is a spooky, mysterious relation; a relation that cannot be explained without begging the question (it is that which necessitates all Fs being Gs in a world where $N(F,G)$ holds?). If there is a natural necessitation relation, it is a primitive, incomprehensible entity. $R(F,G)$ on the other hand, is easy to grasp in terms of 'impressions' that even Hume would admit, and yet it still allows us to assign a high probability to our inductive inferences being correct. So what is it about $N(F,G)$ that makes it a better explanation for all ravens having been observed to be black than $R(F,G)$? It seems to me that $R(F,G)$, if by nothing other than parsimony, comes out as the better explanation.

Conclusion

Armstrong argues that the regularity theorist's pattern of inference to the properties of future observables is nothing more than $e \rightarrow h$ (observed to unobserved). He argues that we need an *explanation* of why there is a constant conjunction between certain universals, and that natural necessitation is the only option.

However, I have shown that what is required is not an explanation of why (say) ravens are black, but an explanation of why we observe them to be black, and of course "all ravens are black" serves as a perfectly good explanation for why every time we observe one, it appears black. Furthermore, I have argued that the pattern of reasoning to the universally quantified conditionals that make

¹⁶ Armstrong would claim he has 'inference to the best explanation' to justify inferring $N(F,G)$, but given that I deny natural necessitation relations to be the best explanation for regularities, this response holds little weight.

Regularity Theory and Inductive Scepticism

up the law is perfectly rational. One doesn't simply move from straight from the observed to the unobserved, but via $R(F,G)$ ¹⁷. The pattern of inference is therefore as follows:

$$e \rightarrow R(F,G) \rightarrow h$$
¹⁸

This pattern of inference *is* rational, and demonstrates how the regularity theorist is in no way committed to inductive scepticism.

University of Nottingham
Nottingham, United Kingdom

¹⁷ Where $R(F,G)$ is the contingent relation between universals derivable when a universally quantified conditional is justifiably inferred via the law of large numbers.

¹⁸ Again, I feel that a Campbell/Franklin-style response gives us more reason to suppose the universally quantified conditional holds than any other distribution of colours (that is, something like "most of the ravens are black, but three are white"), but even if we merely accept that "most ravens are black", we still have good reason to make inductive inferences about the colour of the next raven we observe.

LYCEUM

Bibliography

- D.M. Armstrong. *What is a Law of Nature*. 1983 CUP.
- A.J. Ayer. *Hume*. 1980 OUP.
- M.B. Brown. "Review of Stove 1986" *History and Philosophy of Logic*. 1987 **8**, 116-120.
- S. Campbell and J. Franklin. "Randomness and the Justification of Induction" *Synthese*. 2004 **138**, 79-99.
- P. Castell. A Consistent Restriction of the Principle of Indifference, *British Journal of The Philosophy of Science*. **49** (1998), 387-395.
- D. Hume. *A Treatise of Human Nature*. 1985 Penguin Classics.
- D. Hume. *An Enquiry Concerning Human Understanding*. 1999 OUP.
- B. Indurkha. "Some Remarks on the Rationality of Induction" *Synthese*. **85**(1), 95-114.
- C. Hempel. *Aspects of Scientific Explanation*. 1965 Englewood Cliffs.
- D. Lewis. *Philosophical Papers*. Introduction 1986 OUP.
- P. Lipton. *Inference to the Best Explanation*. 2004 Routledge.
- K. Popper. *Conjectures and Refutations*. 1969 London: Routledge and Kegan Paul.
- D.C. Stove. *The Rationality of Induction*. 1986 Clarendon Press.