Matters of Metaphysics

D. H. MELLOR
Professor of Philosophy, University of Cambridge

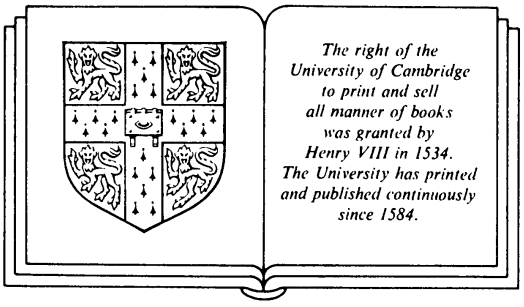

CAMBRIDGE UNIVERSITY PRESS
Cambridge
New York Port Chester
Melbourne Sydney

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 2RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521411172

First published 1991

A catalogue record for this publication is available from the British Library

ISBN-13 978-0-521-41117-2 hardback
ISBN-10 0-521-41117-3 hardback

Transferred to digital printing 2005

15 *The warrant of induction*

1 Introduction This lecture will last less than twenty-four hours. I know that, and so do you. And you knew it before I said so. How? Because you knew that lectures don't last twenty-four hours. How do you know that? You haven't heard this one, and 'for all you know' (as the saying is) I could go on all night. But you know I won't. And the 'all you know' which tells you that, without entailing it, is the fact that none, or almost none, of the many lectures, on all subjects, which you've heard or heard of, have lasted that long. If many of them had, you wouldn't have known that this one won't; but as it is, you do know that.

That's a piece of induction. We believe that something (a lecture) has an observable but as yet unobserved property (a short duration), and this belief is warranted by the fact that many otherwise diverse things of the same kind (lectures) have all – or nearly all – been observed to have that property. The problem, set for us by David Hume in 1739, is to say why such observations warrant such a belief.

The greatest of all Cambridge philosophers, Frank Ramsey, said in 1926 that 'we are all convinced by inductive arguments, and our conviction is reasonable because the world is so constituted that inductive arguments lead on the whole to true opinions' (Ramsey 1926 p. 93). As usual, he was right, but too brisk for most philosophers. Most are still unpersuaded, despite the subsequent work of Professor Braithwaite (1953 chs 6–8) and others. Tonight I will try to say, less briskly, but still within twenty-four hours, why I think Ramsey was right.[1]

First I must say what it is to warrant a belief. Or rather, what it is about a belief that is to be warranted. What's to be warranted is its truth, that is the truth of what's believed: for example, that this lecture will last less than twenty-four hours. That's why observations which warrant a belief can't also warrant a contradictory one: contradictory beliefs can't both be true. Yet our

[1]What follows of course also owes much to other writings on induction and knowledge, notably Ramsey (1929b), Armstrong (1973), Goldman (1980), Grandy (1980), Hacking (1980), Levi (1980b), Nozick (1981 ch. 3) and van Cleve (1984).

inductive warrant doesn't entail that the belief it warrants is true. The belief that this lecture will stop tonight would be warranted now even if I went on to make that belief false; and the belief I would then make true would still not be warranted now. But how can a warrant be valid and yet fail? That's part of the problem of induction.

The problem isn't just that your inductive evidence fails to entail the belief it warrants. It's worse than that. For that belief is a prediction, about when this lecture will end. And its success as a prediction will be settled in due course by an observation: for example, by your consulting a watch when the lecture ends. Some such observation will settle that prediction regardless of its present inductive warrant. If I keep you here till dawn, your present inductive evidence that I won't will count for nothing against the evidence of your senses that I have done. And if I stop sooner, your senses will need no help from induction to tell you so. You know that your inductive warrant for what you now believe is negligible compared with the warrant that some future observation will give it – or deny it.

Yet we trust such predictive beliefs incessantly, often with our lives. Take the belief that this building will stay up while I speak two more sentences. We have, I hope, strong inductive warrant for that belief, and thus for staying here at least that long. Yet the belief is still only a prediction (comma), because we know that its present warrant will be quite superseded by that of our senses ... now. But how can a warrant that is now so weak ever have been strong enough to trust our lives to? And as in this case, so in general: how can a prediction ever be better than a guess? That's another way of putting the problem of induction.

2 Induction and observation I'm taking induction to be something that can warrant anticipating an observation, the problem being to say how. This reading of induction sets limits to the problem, and makes assumptions about observation, which I must briefly defend. It limits the problem by excluding other cases in which evidence might warrant a belief that it doesn't entail. Take mathematics. All even numbers so far checked are the sum of two prime numbers. This doesn't entail that they all are. Does it warrant that belief? I don't know; and if it does, I don't know how: but not, I suppose, in the way that beliefs about how long this lecture will last are warranted. So no one should object to my account of that, just because it won't also explain the force (if any) of inconclusive evidence in mathematics.

More seriously, my account won't explain everything that has been called induction in science. It won't explain how observations give rise to new

theories: because new theories typically predict new kinds of entities (new subatomic particles, say), which what I call induction doesn't. Nor will it always explain how observations warrant existing theories. A theory's entities may not (yet) be observable; and until they are, induction can only warrant what the theory implies about other things that are observable. How that can warrant believing the theory's unobservable implications – if it can – must be quite another matter.

Induction may not be enough to warrant a theory in the sense of warranting belief in it or in something like it. But it is necessary. Since warranting a belief means warranting its truth, it also means warranting anything it entails. And theories entail generalisations which if true state laws about the observable phenomena they explain. To warrant a theory we must warrant those laws. But we can't warrant them by observation directly, since we can only ever observe instances of a law, never the law itself. For laws may have future instances as well as past ones, and since we can't observe the future, we can never observe that a generalisation will never be refuted by some future instance: some interminable lecture, or some wayward subatomic particle. So we can warrant a law by observation only by observing some of its past instances, and then only if this warrants predicting, of any other instance, that observing it would give the same result. But this is what I call induction. Induction may not be the whole story about how observation can warrant believing laws and theories; but it must be a part of it.

Now to observation. I've said that observation is what settles predictions, including theoretical predictions. This doesn't mean that I think observation is independent of theory. We must hold some microphysical theory in order to see a bubble chamber track as the track of a subatomic particle. But a prediction about that particle's track may still be warranted inductively by observations of previous tracks, and be refuted by observing this one. The fact that observers need theories doesn't stop theoretical predictions being warranted inductively by past theoretical observations, or overthrown by future ones.

Nor need observations be infallible to settle predictions. Of course observations can yield unwarranted beliefs (in mirages, for instance); they can need inductive support; and their warrant can sometimes be outweighed by inductive evidence. A mere glimpse may not warrant the belief that there are sparrows in the park, unless sparrows have been seen there many times before. It may take a long look by a good birdwatcher to warrant believing in the presence of a bird that's both nondescript and rare. But all this shows is that inductive evidence can outweigh a glimpse, not that it can outweigh

observation altogether. It can't, since inductive evidence consists of observations of the very kind that it lets us anticipate: observations of birds, subatomic particles, lectures. Seeing that some birds are sparrows can't warrant predicting that a new one is as strongly as it warrants the belief (that the old ones are) on which this prediction is based. But then, by the same token, an equally good look at the new bird must warrant a belief about it more strongly than the observations of those other birds could: so we must be able to settle the prediction that it's a sparrow by looking at it. In short, observation can always outweigh an inductive warrant for any prediction about anything observable – as everyone knows who's ever been unpredictably stood up.

But observations are still fallible, a fact we obscure when we use words like 'observe' so that by mere definition we can't observe what isn't there. So I will not use 'observe', 'see', 'hear', etc. in that misleading way. When my looking to see if there are sparrows makes me believe that there are, I shall say that I see that there are, whether there are or not. For there might not be: our senses always can give us false beliefs and sometimes do. Observation no more entails the truth of the beliefs it gives us than induction does. How then can it warrant them?

Well, not of course by induction, so answering this question is no part of my present task. But answering it will help, by enabling me to reset a standard of warranted belief which has hitherto made the problem of induction seem insoluble. For as we have just seen, observation must be able to warrant any beliefs that induction can warrant. Yet by the standard that induction is generally condemned for failing, observation couldn't warrant them. But that's incredible: no one really doubts that observation can and does warrant most of the particular everyday and scientific beliefs it gives us. So we need a more credible standard of warranted belief anyway, one that observation can meet. If we can find such a standard, and then show that, and how, induction can meet it too, the problem of induction will be solved. And that's how I propose to solve it.

3 Observation and warranted belief How should a standard of warranted belief be set? A standard implies some goal or object whose achievement meeting the standard should ensure. Take the safety of goods, which is the object of many British Standards. Meeting those Standards should ensure that goods are safe: that is the test of a good safety standard. Here the object is not safety but truth. To warrant a belief is to warrant its truth. So the test of a

good standard of warranted belief is this: would the truth of a belief whose warrant meets that standard be thereby ensured?

How can a warrant, like that of our senses, ensure the truth of a belief it doesn't entail? Well, suppose I see that there are sparrows in the park. This belief will be true if and only if there really are sparrows in the park. So my observation should warrant my belief if it's such that, if I'm seeing sparrows there, then they're there.

What would ensure the truth of this conditional proposition? We know an entailment would. If my seeing sparrows entailed that they were there, the proposition that they're there if I'm seeing them would be not merely true but necessarily true. And if that were so, then seeing them would certainly warrant believing they were there. But it isn't so. Seeing sparrows doesn't entail that they're there. Our conditional proposition is not a necessary truth.

Still, propositions can be true without being necessarily true. Suppose you have a mass of M units. This fact makes true a whole raft of conditional propositions of the form 'if a force of F units is applied to you, you will accelerate at F/M units'. And since your mass is contingent – we could all be more or less massive than we are – those conditionals are all only contingently true.

Now suppose some similar contingency makes it true that if I'm seeing sparrows then they're there. Why can't such a contingency make this observation warrant that belief? After all, it still can't also warrant the contradictory belief that sparrows aren't there. The conditionals 'they're there if I'm seeing them' and 'they're not there if I'm seeing them' can't both be true while I'm seeing them, however contingent they are. Nor could the warrant both hold and fail. The sparrows can't fail to be there when I'm seeing them while the conditional 'they're there if I'm seeing them' is even contingently true.

So far so good. But what might our contingent warrant be? Not just that sparrows happen to be there while I'm seeing them. That does arguably make it true then that they're there if I'm seeing them. But if so, that conditional is then only made true by the truth of the very belief it's meant to warrant: namely, the belief that the sparrows are there. So if that were the warrant, the truth of that belief would in effect be warranting itself; and it can't do that.

For our purposes we need something more, to make a conditional of the form 'if P, then Q' true, than the mere fact that P and Q are true: something more like mass. Your mass doesn't only entail an infinity of conditionals, about how you would accelerate under various forces. It guarantees that, while you keep it, the same force would always give you the same acceleration. Mass is a deterministic disposition: if a force F is applied to a mass M, the chance –

the physical probability – that it will accelerate at a rate F/M is 1. And that makes the conditionals which a mass entails non-trivially true: since the mere fact that M does accelerate at a certain rate doesn't entail that it had no chance of doing otherwise.

Your mass in fact embodies a causal link between forces and accelerations: it's what makes forces cause you to accelerate at a certain rate. And just such causal links, between us and what we observe, are what enable our observations to warrant the beliefs they give us. To see, for example, is to be affected by light which something emits or reflects into our eyes, the relevant effect being that we get a belief about the source of that light. I see a sparrow when the light from something causes me to believe that it's a sparrow. My belief is true if it is a sparrow, and warranted if the causal link between me and it makes it non-trivially true that, if I believe it's a sparrow, then it is.

How does the causal link make this conditional proposition true? By making sparrows cause me to believe that there are sparrows there, and other things cause me to believe that there aren't. If the causation is deterministic, the chances of both these effects will be 1, from which it follows non-trivially that if I do see sparrows, then they're there. That's what makes my seeing them warrant my belief that they're there.

The causal links, and hence the chances, which this warrant depends on are of course contingent on many facts about me and the sparrows: that I have the concept of a sparrow, decent eyesight and my eyes open; that the sparrows are not only there but well lit and in view; and so on. Those facts might have been otherwise: so my belief might not have been warranted. But as things are it is warranted, since as things are there's no chance of my observation giving me a false belief.

But causation needn't be deterministic. An effect's chance must indeed be greater if a cause occurs than (in the circumstances) it would be if it didn't occur (see chapter 13). But its chance with its cause may still be less than 1, and its chance without it may be greater than 0. Smoking may for example cause a man to get cancer even if he isn't certain to get it if he does smoke, and even if he might get it without smoking.

And even an indeterministic causal link can make my looking for sparrows warrant my beliefs about them, by making their presence nearly always cause me to see them, and their absence nearly always cause me not to see them. The warrant is worse, of course, when it makes my chance of getting a true belief less than 1. But a 99% or 95% chance may still be good enough: it will still make my seeing sparrows warrant the belief that they're there far more strongly than the belief that they're not.

This kind of chancy causation moreover enables observation to be fallible, since it lets me sometimes see sparrows that aren't there. Now when that happens, it isn't true that the sparrows are there if I'm seeing them. So that conditional needn't be true after all. Observation can warrant a false belief, provided it gives the belief a small enough chance of being false.

This, I maintain, is how the greater chances that effects have when their causes occur than when they don't enables observation to warrant beliefs: contingently and fallibly, but quite well enough. And chances which can do that can enable induction to do so too. But before showing how, I must briefly dispose of two objections to the standard of warranted belief that I've so far been appealing to.

4 Chance and causation The first objection is to my concepts of chance and causation. To serve our turn, chance must be more than a merely epistemic probability: that is, more than a mere measure of how strongly beliefs are warranted. To say that a belief is warranted when it's epistemically very probable is just to say that it's warranted when it's warranted. This doesn't tell us what the warrant is, nor why it is one. Talk of chance can tell us that only if chance can be specified and detected independently of how it sometimes warrants our beliefs, but yet in a way which shows how and when it does warrant them. And so it can be.

We specify and detect chances mainly as indeterministic causes of statistical facts. Suppose a particular coin tossed in a certain way W has a chance C (between 1 and 0) of landing heads. This entails the existence of a statistical law, namely that if enough sufficiently similar coins were tossed this way, there'd be a very high chance that a fraction close to C would land heads. Suppose many such coins are so tossed, and the fraction that land heads is in fact close to 0.6. The chance of that happening is greater if C is 0.6 than if it isn't. And if C is 0.6, that is the cause: the fact that the chance of each of these many W-type tosses landing heads is 0.6 is what causes – and thereby explains – the fact that the fraction which do land heads is close to 0.6.

Many statistical facts about distributions of physical and biological properties are similarly caused by chances. Take a pure lump of the most common isotope of radium. In any given time a closely constant fraction of its myriad atoms will decay: in 1622 years, about 50%. That's because – literally because – each of those atoms has this chance of decaying in that time. Or take the fact that about half the foetuses conceived in large populations are female. This too is an indeterministic effect, of each of those many conceptions having a 50% chance of producing a foetus with only X-chromosomes. And so on.

Chances also have causes. Bending a coin affects its chance of landing heads when tossed. Bombarding a radium atom raises its chance of decaying. Starting to smoke causes people's chances of getting cancer (and hence the fraction of people who do get cancer) to increase.

Chances, in short, don't only help to link other things as cause to effect. They too are causes and effects: because they have causes and effects, by which we can specify and detect them, quite independently of how they enable observation – and, as we shall see, induction – to warrant our beliefs. Most chances actually warrant no beliefs at all. They only warrant them in observation because there the effects of some fact include someone's getting a belief: ideally the very belief which that fact makes true, like the belief that there are (or that there aren't) sparrows in the park. This is what turns the chance C of those effects into a probability that whichever belief is thereby caused is true. And what then makes C an epistemic probability is its entailing that many such observations would give us a very high chance of getting a fraction of true beliefs very close to C. So if C is 1, such observations would always give us true beliefs. And if it's very close to 1, they would give us a very high chance of almost always getting true beliefs. What better measure could there be of the prospects of truth which observation gives the beliefs it produces, and hence of how strongly it warrants them?

5 Knowing and knowing you know Well, one answer to that rhetorical question is: a self-intimating measure. Surely a warrant for a belief of mine is only worth having if I know I have it, so that I know which of my beliefs are warranted and therefore worth trusting? But this standard for warranted beliefs doesn't make belief-warrants self-intimating. This is the second objection to it.

Now I could just add to my standard the rule that belief-warrants must be self-intimating. But I won't: because that rule in fact makes beliefs almost impossible to warrant, while doing nothing to ensure either that warranted beliefs are true or that their warrants are useful to those who have them.

This all follows from a conceptual link between warranted belief and knowledge. The link is that knowing about something is having a true and warranted belief about it. I will call this the 'true warranted belief' thesis. I assumed it at the start when I said that you knew this lecture would last less than twenty-four hours, and inferred that you had a true and warranted belief to that effect. And although the thesis is disputed, I shall go on assuming it, just because it causes trouble and is at least plausible. So I'd better assume that it holds.

The trouble it causes is this. Suppose, to vary the example, you have perfect eyesight and look in good daylight at a green frog on a red leaf. This causes you to get a belief about the frog's colour which in these circumstances has a very high chance of being true. Suppose it is true: you get the belief that the frog is green. I say your observation warrants this belief (belief 1). So, by the true warranted belief thesis, you know the frog is green.

Now try to make this warrant self-intimating: if you have it, you know you have it. That is, by the true warranted belief thesis, you believe you have it, and this belief (belief 2) is true and warranted. So, since this warrant too must be self-intimating, you believe belief 2 is warranted; and *that* belief (belief 3) is warranted. So you also believe that, that belief (belief 4) is warranted and so on, rapidly *ad infinitum.*

If you can believe all that, you can believe anything. In fact, I doubt if you have any of these extra beliefs, and you certainly won't have all of them. But you don't need any of them. Consider for example that part of the warrant for your belief that the frog is green must be that you aren't red–green colour-blind. But you may well not believe this, because you may never have heard of colour-blindness. That can't stop you knowing that the frog you see is green, or no one could have known anything about colours before the phenomenon of colour-blindness was discovered. So your warrant for believing that the frog is green can't have to be self-intimating: you must be able to know that something is green without knowing that you know that.

If belief-warrants, and therefore knowledge, had to be self-intimating, you could know almost nothing that observation tells you. Suppose you do believe you aren't colour-blind. To be knowledge, this belief must be warranted. How? By your having passed a colour-blindness test? But for this warrant to be self-intimating, you must believe you've passed a colour-blindness test, and that belief must be warranted: say by the test having been tested to see if it really can pick out people who are colour-blind. But how could that have been done if, until it is done, no one can know whether anyone is colour-blind?

The fact is that if, to know something, you had to know you knew it, knowledge would entail a series of warranted beliefs which observation does not, and generally could not, give us. If observation is to warrant even a few of the beliefs it gives us, this thesis, that you always know what you know, will have to go. And so it should: because the thesis sets a quite unwarranted standard for warranted belief. That is, it does nothing to ensure that warranted beliefs are true, which is the whole point of warranting them. My knowing that if I make an observation my chance of getting a true belief is high won't in general make the chance any higher. Most of the beliefs my senses give me

are warranted no more strongly when I know they're warranted than when I don't.

Nor is it true that belief-warrants are worth having only when we know we have them and therefore know which of our beliefs to trust. We can't only trust – that is, be disposed to act on – some of our beliefs and not others. Not to trust a belief is not to have it. Suppose you want a frog, and pick up that green thing on that leaf because you believe it is a frog. If you did stop believing that this belief of yours was warranted, you might indeed no longer trust it in this way – because you would no longer have it: you would no longer really believe that the green thing was a frog.

This doesn't mean that it isn't useful for beliefs to be warranted in my sense. That is useful, because it's useful for beliefs to be true. Truth is the attribute of beliefs which ensures that the actions they cause will succeed.[2] For example, if your belief, that what you're picking up is a frog, is true, your action will succeed: it will get you what you want, namely a frog. And if this belief is warranted, in my sense, it will have a high chance of being true: so your action will have a high chance of success. This is what makes it worth having warranted beliefs. Your knowing that they're warranted would only make them more worth having if it raised their chances of being true. And generally it doesn't.

In short, the standard for warranted belief set by the 'we know what we know' thesis is both unattainable and ineffective. It's no objection to observation or induction that they can't meet it. Yet, as we shall see, this is the objection that has made the problem of induction seem insoluble. For the 'we know what we know' thesis is still oddly attractive, despite its absurd consequences. Why it's so attractive, I'm not sure. I suspect it's because we still think of ourselves as essentially conscious, and especially self-conscious: so we like to think that if we know something, we know we know it. But the fact is that we needn't, and we usually don't, know what we know. And once we recognise this fact, we can solve the problem of induction.

6 Induction Let me call events which last less than twenty-four hours 'terse'. This event, as you see, is a lecture. It's also terse, but that you can't yet see. You predict it, by inferring it: this is a lecture, so it's terse. And this inference displays a general disposition or habit: you would have drawn it of any lecture, not just this one. And this inferential disposition, like your mass (your inertial disposition), embodies a causal link: just as your mass makes forces

[2]See chapter 2.4 and Whyte (1990).

cause you to accelerate, so your inferential disposition makes your coming to believe that an event is a lecture cause you to predict that it's terse.

We've already seen how causal links can enable observations to warrant beliefs: namely, when the fact that makes a belief true is what causes you to get that belief. The rôle of causation in an inference is different. What causes you to believe the conclusion of an inference isn't the fact that makes it true, but the fact that you've come to believe the premise of the inference. How can this warrant the conclusion?

Suppose, to simplify the discussion (it's not essential), that your inferential habits are deterministic, like your mass: this habit would always make you infer that a lecture was terse, never that it wasn't. And suppose that every lecture has some chance of being terse. Then whenever your premise ('this is a lecture') is true, your conclusion ('this is terse') has some chance of being true. And if this chance is high enough, your prediction is warranted.

I shall call an inferential habit warranted (or good) if, whenever the premise is true, the conclusion has a high enough chance of being true. For then there's a very high chance that all or nearly all the conclusions of many such inferences will be true when their premises are. A particular inference may still of course fail to warrant its conclusion, because its premise may be false.[3] But that's no fault of the inference. All an inference can do is give a conclusion as high a chance of being warranted as its premise has of being true. That's good enough to warrant calling it warranted, or good.

We have of course many habits of inference. Some perhaps we're born with; most we acquire. And, as Hume remarked, we acquire most of them by induction. Mostly, the more we see that otherwise diverse things with one property F (being a lecture) all or nearly all have another property G (being terse), the more we tend to predict that other things which are F are also G.

This isn't a universal tendency: it doesn't apply to all observable properties.[4] Take ageing. The more years we adults do survive, the more we tend to predict that we won't survive another one. We recognise many such 'counter-inductive' phenomena: metal fatigue and caterpillars turning into

---

[3]This stops counter-examples to the 'true warranted belief' thesis which include inferences from premises that are warranted but untrue, as in Gettier (1963). But it doesn't stop such an inference being part of a two-stage causal process of observation and inference that still lets the observation warrant the conclusion by giving it a high chance of being true.

[4]I don't count Goodman's (1965) 'grue' and its many progeny. Green, like blue, is an observable property: so what 'grue' denotes isn't. When I see truly that a frog is green (and hence grue) my belief is caused by the fact that it's green, and not at all by the fact that I see it less than so many years after the birth of Christ (see Blackburn 1969, Jackson 1975). How we know green is a property is another question, which I needn't answer: it suffices here that it is, because belief-warrants needn't be self-intimating. Disposing of the 'we know what we know' thesis also disposes of Goodman's 'new riddle of induction'.

butterflies are two more obvious examples. But we recognise them inductively. Our mortality statistics show that the observed chances of adults surviving another year have so far always or nearly always decreased with age. That's why we tend to predict that ours will. And similarly in the other cases. Our counter-inductive tendencies always have an inductive basis.

The question is: what warrants the inductive basis? We certainly think it's warranted. Counter-inductivists aren't just odd: they're mad. Imagine one. He won't eat bread: he thinks it would poison him, because it's never poisoned anyone before. He would eat cyanide, which he also expects to freeze in the oven and bake in the fridge; but not by swallowing it. He won't use any language people have so far understood, or breathe air, or drink water. And throughout his (brief) life he consistently defends his wholesale counter-inductivism by predicting that as it's almost never worked yet, it will now.

And so it could. He could be right. He could outlive us all. But he won't; and we know he won't. How? What makes induction a better basic tendency than counter-induction?

Take any pair of basic properties which we've seen to be correlated, like being a frog and being green. Suppose all or nearly all the many frogs so far seen have been green. The more, and the more diverse, those frogs and their surroundings, the more we tend to predict that other frogs will be green. Exactly how many, and how diverse, doesn't matter: what matters is what warrants this general inductive tendency.

The tendency increases both with the number of frogs seen, and with their diversity. Consider first the way it increases with the number. So suppose the frogs (and their surroundings) are all of a kind: tree frogs, for example. The more tree frogs are all seen to be green, the more we tend to predict that other tree frogs are green – and the more counter-inductivists tend to predict that they're not. What makes ours the better tendency?

That depends on what chance a tree frog actually has of being green. It may be a law of nature that all tree frogs are green: that is, the chance of any tree frog being green may be 1. Suppose it is. Then all observed tree frogs will be green, so induction will always make us predict that other tree frogs are green. And this inference couldn't be better: whenever its premise ('this is a tree frog') is true, the chance of its conclusion ('this is green') being true is 1. Whereas counter-inductivists, seeing only green tree frogs, will always predict that other tree frogs aren't green: an inference which couldn't be worse.

So much the worse for counter-induction. And it's no better off if the law is that no tree frogs are green. For then no observed tree frogs will be green,

induction will always yield the good habit of inferring that others aren't either, and counter-induction the bad habit of inferring that they are. And as for frogs and colour, so for all basic observable properties. Whenever they're linked by a deterministic law, induction will always yield good habits of inference, and counter-induction will always yield bad ones. In all such cases our inductive tendency is warranted, and the counter-inductive one isn't.

So far so good – provided these warrants needn't be self-intimating. I say the law that all tree frogs are green warrants my inferring that something is green from the fact that it's a tree frog. But suppose I must know that I have this warrant. Then I must know this law. So I must believe it, and this belief must be warranted. But the law entails the very inference which it's meant to warrant: tree frogs can't all be green unless this one is. So unless my inference is warranted already, my belief in the law won't be warranted. Thus to claim that the law is what warrants this particular application of it simply begs the question of whether it's warranted at all.

This is the stock objection to contingent solutions to the problem of induction: they beg the question. And so they would if belief-warrants had to be self-intimating. But as we've seen, they don't. The law that all tree frogs are green can warrant the habit of inference which induction will then give me, just because I needn't know that it does. I can know by induction that a frog is green without knowing that law, just as I can know that it's green by looking at it without knowing I'm not colour-blind. I may know the law, just as I may know that I'm not colour-blind; but I needn't. So my saying that the law is what warrants this induction doesn't beg the question.

Deterministic laws warrant induction. And so do statistical laws, and hence the chances that entail them. For suppose it's a law that all tree frogs have a certain chance of being green. The greater this chance, the better the inference that a tree frog is green. But also the greater the chance that many tree frogs will all or nearly all be green. In particular, the greater the chance that all or nearly all observed tree frogs will be green – and hence that we inductivists will infer that others are too. In short, the better the inference, the more likely we inductivists are to make it. And the less likely counter-inductivists are (since the better it is, the less chance there is of all or nearly all observed tree frogs not being green). So far still so good.[5]

---

[5]If tree frogs' chances of being green are all 50%, the inferences, that this one is green, and that it's not, are equally unwarranted. So therefore is induction, which has the same (minute) chance of making us make either inference as counter-induction has. So induction fails here, but only because there's no success to be had. There will of course be a good inductive inference to the fact that this tree frog's chance of being green is about 50%, but that's another matter.

Furthermore, the more tree frogs we see, the less risk we inductivists run of making inferences that aren't good. For when a tree frog's chance of being green isn't close to 1, the more tree frogs we see, the less the chance that they'll all or nearly all be green and so induce us to infer – badly – that others are. That's what warrants the way in which our tendency to infer that tree frogs are green gets stronger as we see more and more tree frogs which are all or nearly all green.

That's how induction is warranted (and counter-induction unwarranted) by simple laws linking a pair of observable properties. But laws are rarely so simple. A frog's chance of being green may depend on many things. Hence the other factor in induction: the way in which our tendency to infer that frogs are green also increases with the diversity of the frogs which we've seen to be all or nearly all green.

Suppose you and I both see many frogs, but I only see tree frogs, while you also see many frogs of other kinds. How will this help you to make better inductive inferences about frogs? It won't if all frogs have the same chance of being green (or whatever): then we're on a par. But it will if they don't. For suppose tree frogs have a high chance of being green, and others don't. Then the inference 'it's a frog, so it's green' is good when it's a tree frog and bad when it's not. And because my chance of seeing only green frogs is now greater than yours (since I only see tree frogs), my chance of being induced to make this inference when it's bad is greater than yours. Whereas, provided you notice that tree frogs differ from other frogs, your chance of seeing only green tree frogs, and hence of being induced to make the good inference 'it's a tree frog, so it's green', is just as great as mine.

In short, whenever diversity in our inductive data matters, it's better to have it than not, and better to notice it than not. Of course we won't know if it matters at the time, since that will depend on laws we don't yet know. So the only warrantable tendency we can have is to be always readier to make inductive inferences when we know our data are diverse than when we know they aren't. And this is the tendency we do have.

But what if our unseen frog has no chance, high or low, of being green? In other words, no law, deterministic or statistical, simple or complex, links being a frog of whatever kind this one is with being green. Then there's no good inference from this being a frog of any such kind to its being green, or to its not being green. Neither of the inferential habits which induction or counter-induction might give us is good. But equally there's no specific chance, high or low, of induction or counter-induction giving us either habit, since there's no specific chance of many frogs of such a kind being all or

nearly all green, or all or nearly all not green. In this lawless situation, neither induction nor counter-induction is warranted.

How do we know this isn't our situation? Maybe we don't know. But we don't need to know. Induction only fails here because here there's no success – no good habit of inference – to be had. When there is a good habit to be had, induction will always give us our best chance of getting it; and the better the inference, the better the chance. And that's enough to warrant induction as a general basic tendency.

But is it enough for induction to warrant the specific inference that this lecture is terse? Well, it is if lectures (or at least Inaugurals in English) have a high chance of being terse. For this will both warrant the habit of inferring that such lectures are terse and give induction a high chance of inducing this warranted habit, as it has done.

But do lectures like this have a high chance of being terse? Of course they do. You may not know this (though I dare say you do), but that doesn't matter. All that matters is that they do. And this fact (which, since you do at least believe it, you can't honestly deny) entails that your belief that this lecture is terse is both warranted, and warranted by induction. And so you do know already that it's terse – that it will last less than twenty-four hours – because, as you'll see at the end of this sentence, this belief of yours is not only warranted, it's true.