Chapter 4. The Problem of Induction: One Problem or Many?

It is standard in philosophy to divide the pathways to knowledge and justification into five broad categories: perception, a priori reflection, memory, testimony, and inductive reasoning.[1] Each one of these categories is usually given separate treatment since they each raise their own distinctive philosophical problems.

All of these categories are indispensable to a working system of knowledge. If all I had was perceptual knowledge, then all I could ever know would concern my immediate vicinity in the current moment. A priori reflection gives me knowledge of mathematics, logic, and philosophy, but it doesn't tell me anything about the contingent reality that I inhabit (with one exception—that *I exist*). With memory, I get to extend my knowledge to matters that concern the vicinities that I've inhabited in the past, but this doesn't give me any knowledge about matters that I haven't personally observed. With testimony, I get to learn about the matters that other people have observed.

But all of this is still far away from the full body of knowledge that we all think we have. I think that I know that all cats weigh less than a thousand kilograms. But neither I nor anyone else has observed all of the cats that there are. It is consistent with all of my observations that there is a thousand-kilogram cat living somewhere out of sight. I also think that I know that all water boils at a hundred degrees celsius at sea level. But that said, nobody has ever tried boiling *all* of the water in the universe to check.

Much of what we think of as scientific knowledge consists of generalizations that extend beyond our finite observations. Scientific laws are stated as universals: e.g. that *all* light travels at 3 × 108 m/s, that *all* samples of water boil at a hundred degrees celsius (at sea level), that *all* cats are relatively medium-sized compared to other land mammals. Yet, nobody has observed all instances of light travelling, all samples of water, or all the cats. In each case, scientific knowledge depends on generalization. And yet, the universal form of these laws is crucial because, without it, we wouldn't be able to use these laws to make future predictions. It is because I know that *all* water boils at a hundred degrees celsius (at sea level) that I now know that the tea I'm brewing is about a hundred degrees. I didn't have to check using a thermometer.

Since we haven't observed all of the relevant cases, how do we gain knowledge (or justified belief) of these laws? The standard answer is that we use *inductive*

[1] There are numerous examples of philosophers who attempt to subsume one or more of these categories into another. But let's ignore that. In my view, they should be thought of as distinct.

reasoning. Inductive reasoning is supposed to take us from instances of things we've observed to knowledge, or justified belief, about things that we haven't observed.

But what is inductive reasoning, exactly? We usually define induction by contrasting it with deduction. Briefly put, deductive reasoning is the kind of reasoning or argument whereby we attempt to establish the conclusion on the basis of the premises as a matter of logical necessity. When a deductive argument is successful—that it, it is valid—then its conclusion is guaranteed by the premises. The conclusion cannot be false unless one (or more) of the premises is false.

Induction, on the other hand, is supposed to be a mode of reasoning, or a kind of argument, where the conclusion is not strictly *guaranteed* by the premises, but is nonetheless rendered probable by the premises. Consider, for example, all of my observations of house cats. Imagine that every time I've seen a cat, I've taken a record of its weight. I thereby amass a long list of premises:

1 Snuggles the cat weighs less than a thousand kilograms.

2 Fernando the cat weighs less than a thousand kilograms.

3 Oscar the cat weighs less than a thousand kilograms.

...

N Mr. Sam the cat weighs less than a thousand kilograms.[2]

Let's say that premises 1 through N record every cat I've ever seen and I've never seen one weigh anywhere near a thousand kilograms. Given all of these premises, it seems that I'm entitled to conclude:

C All cats weigh less than a thousand kilograms.

The important thing to notice is that, although the conclusion seems reasonable given my observations, it doesn't follow from the premises as a matter of logical necessity. It is possible for me to make all of these observations, and thereby know the truth of all of these premises, when in fact the conclusion is false. Perhaps there is some thousand-kilogram cat hiding out in a forest somewhere.

Another way to see that the argument isn't strictly valid is to compare it with another argument with the same form. Suppose that I happen to have seen the same number of ravens as I've seen cats, and every time that I've seen a raven, I have seen that it is black. I thereby construct the parallel argument:

$1'$ Raven # 1 that I have observed is black

$2'$ Raven # 2 that I have observed is black

$3'$ Raven # 3 that I have observed is black.

[2] Yes, these are real cats.

...

N′ Raven # N that I have observed is black.

C All ravens are black.

The structure of this argument is a perfect mirror of the last one. Not only that, but it is even plausible to think that the conclusion is justified. Nonetheless, the conclusion isn't true. Even though they are rare, there do exist white ravens, and so it's not true that *all* ravens are black. But this just goes to illustrate what we already know about induction: that even when an inductive argument appears to be good, the truth of its conclusion isn't guaranteed.

The reason that I would like to devote an entire chapter to induction is that there is a longstanding tradition (particularly within Western philosophy) of singling it out as an especially problematic source of knowledge and belief. The claim is that we can never obtain justified beliefs or knowledge on the basis of induction, even if we can obtain justified beliefs and knowledge from every other source. Induction is supposed to be *uniquely bad*, according to this line of thinking. Let's call this position "induction skepticism." The source of this brand of skepticism is David Hume and his notorious problem of induction.

It is my view that the so-called problem of induction has been highly overrated ever since Hume made it popular. Indeed, it is not even very clear what the problem is supposed to be. There are, to be sure, various formulations of induction-specific skeptical arguments that have cropped up as homages to Hume, and we will get to them momentarily. But there doesn't seem to me to be a single common worry that unites each version. Nonetheless, I will respond to each problem of induction as they appear.

In saying that the problem of induction is overrated, let me be clear as to what I am and am not claiming in this chapter. I am *not* saying that there aren't difficult questions as to *how* we should reason from particular observations to general conclusions. Indeed, there is a huge body of work in the philosophy of science as to how inductive procedures, methodologies, and heuristics work in practice, and how they may be optimized. But it is not my intention to explore those questions here. I'm less concerned about *how* induction works than I am about *whether* it can work at all. In other words, my interest is in the fundamental epistemological question of whether it is possible to have justified beliefs or knowledge on the basis of inductive reasoning. My only opponent is the skeptic who claims that it isn't possible. I will not be substantively engaging in the disputes between friends of induction who argue over the best methods. Perhaps this means that this chapter will be rather unambitious. But then again, there is some value in squaring off with the induction skeptic solely for the fact of their outsized popularity.

Having said that, it is worth being a bit more explicit about the forms of inductive arguments that we will be considering. Induction is usually defined negatively, as it was above; we say that induction is a *non-deductive* form of inference. But that is not very specific. This characterization allows induction, as a general kind of argument, to be a highly heterogeneous and motley assortment. (And indeed, one of the main points that I will make is that the class of inductive arguments *is* highly heterogeneous. It is for this reason that it is doubtful that there can be a one-size-fits-all skeptical argument against induction. But let's not get ahead of ourselves.) Nonetheless, to simplify things, let's focus only on the narrow case of *inductive generalization.* Inductive generalization is the kind of inductive argument that encompasses the two arguments given above. It is the kind of reasoning that iterates premises gained from particular observations and then concludes with a universal generalization. The form of an inductive generalization is always this:

P All observed Ks are P

C All Ks are P.

Again, there are other forms of inductive arguments, but let's ignore them for this chapter. The point is to address the skeptical arguments, and if the problem of induction fails for inductive generalization, then it fails altogether.

Before we consider the original problem of induction from Hume, there is one final clarification to make. So far, I have been ambivalent as to whether we're interested in *knowledge* from induction or *justified belief* from induction. In the broader context of this project, my focus has been on knowledge rather than justification. However, the various versions of the problem of induction are more often focused on justification than knowledge. This makes them rather more extreme than the skeptical arguments of chapter two since they're claiming that our inductive beliefs aren't even *justified.* (Of course, if they aren't justified, then they won't count as knowledge either.) Since my opponents are concerned with justification, I'll meet them at their level and focus on justification as well. My aim, then, is to defend the claim that inductive beliefs may be *justified.* If I'm right about this, then that still leaves unanswered the question of whether they can ever count as knowledge. I think that they can, but I won't press the point here. Suffice to say that the relationship between knowledge and justification is rather complicated, and it would take me too far afield to explore it in this project.

1 The Hume-Russell Problem of Induction

Suppose that I were to attempt an inductive inference in the presence of an induction skeptic. Let's say, as before, that I infer that *all house cats are middling-sized* from the premise that *all of the house cats that I've observed have been middling-sized.* How might the skeptic find fault with this reasoning? How might they argue that my conclusion is unjustified?

Here is a fairly standard way to present the skeptic's point. By making this inference, I must be relying on several unspoken assumptions about how the world works. In this case, I must assume that the sample of cats that I've seen are fairly representative (that there aren't any giant cats inhabiting unknown planets), that, on the whole, cats are fairly similar with respect to their size, that the size of each cat will remain stable into the future (they aren't going to dramatically increase in size overnight), and so on. For if any of these assumptions failed, then my observations would not lend any credence to my conclusion. So, to sum up, my argument is cogent only if we assume that nature is fairly uniform in all of the respects that are relevant to the size of cats.

Similarly, whenever I make any inductive argument, I must likewise assume a certain uniformity in the world or nature. Bertrand Russell, another famous proponent of the problem of induction, would challenge us to justify our belief that the sun will rise tomorrow. We might try to meet his challenge by saying that we have a long track record of observing the sun rising each morning. Russell would reply that these observations only lend support to our conclusion *if* we assume a certain uniformity. Specifically, he argues that we must presuppose that *the future will resemble the past.* For if the future doesn't resemble the past, then there's no guarantee that the sun will rise tomorrow!

Once we admit that all inductive arguments rely on these tacit assumptions of uniformity, our attention should turn to those assumptions. The induction skeptic will want to press us to justify the claim that these underlying assumptions are true. They will say that if we can't justify the implicit assumptions of our inductive inferences, then our inductive inferences won't justify their conclusions. After all, what good would an inductive inference be if the universe wasn't stable, if it was totally random? To put it another way, they take the implicit assumptions for inductive arguments to be *presuppositions* in the sense given in §1.3. According to the definition I gave earlier, a belief B *presupposes* a proposition P just in case justifying P is necessary for being justified to believe B. In this case, the belief in question is the conclusion of an inductive inference (that all cats are middling-sized, that the sun will rise tomorrow) and the presupposition is that nature exhibits the right sort of uniformity.

According to the traditional induction skeptics, we can capture the underlying assumption that all inductive arguments rely on in the form of a single principle. Hume famously called it the "principle of the uniformity of nature." The idea here is that any inductive argument whatsoever—no matter the conclusion and no matter the subject matter—must presuppose that the entirety of nature is ultimately governed by uniform laws. There must be certain fundamental *regularities* that underlie all of nature's course. The universe cannot be *random* or *arbitrary.* For if it were, then there would be nothing stopping the next cat we see from being gigantic, or the earth ceasing to rotate, or the sun popping out of existence.

The principle of the uniformity of nature, which all inductive arguments allegedly presuppose, is notoriously difficult to make precise. The rough idea behind this principle is that, on the whole, nature is stable and coherent. But how do we make this idea any less vague? Hume defined the principle as the claim that *unobserved causes and effects will resemble observed causes and effects.* This will impose a certain sort of uniformity on the world, so that we could expect that all of the cats, who all obey the same laws of cause-and-effect, to resemble each other in size. Russell later simplified the principle to the claim that *the future will resemble the past.* This is supposed to imply, specifically, that the cats we observe in the future will resemble the cats that we've observed in the past.

Now although each of these statements of the principle are vague, let's consider the skeptic's claim that there is such a principle that is presupposed by all inductive arguments. This creates a demand on the proponents of induction to justify this very principle.

It is at this point that the induction skeptics claim to have identified a problem. The problem, they say, is that we *cannot* justify the principle of the uniformity of nature. For suppose that we were to try to justify it. How could we possibly do that?

The skeptics are quick to point out that we cannot *prove* it to be true in the same way that we can prove, say, that every positive integer has a unique prime factorization. We can't demonstrate it to be true by a matter of necessity. After all, it seems possible for the universe to be more random than we think it is. And if it can be more random, then that suggests that we cannot prove that it isn't.[3]

So it appears that we cannot prove *a priori* that the universe abides by the principle of the uniformity of nature. But if that's the case, then it would seem that our only hope left for justifying the principle is to muster some sort of inductive argument. Perhaps we can say that we *observe* the universe unfolding in an orderly, lawlike, uniform fashion. We may then try to conclude, on this basis, that the universe must be law-abiding and uniform.

But if this is our supposed justification for the principle, then the skeptic has caught us in their trap. They have already claimed that *all* inductive arguments presuppose the principle of the uniformity of nature. So if we try to justify this principle by *using induction*, then we are effectively presupposing what we have set out to justify (so they say). We are reasoning in a circle, and circular reasoning cannot amount to justification.

Russell made this point especially poignant by inviting us to justify our belief that *the future will resemble the past.* Again, we can't prove that the future will resemble the past by a matter of logical necessity; the best we can do is extrapolate from our past experiences. But clearly it would be wrong to infer that the future will resemble the past on the grounds that we had observed *past*-futures resembling their pasts. The whole point was to argue that the principle of past-future resemblance will hold *in the future.* So if we appeal to what has

[3]In fact, something like the principle of the uniformity of nature has come up for intense scientific debate during at least one point in our history. During the nineteenth century, before Charles Darwin published *On the Origin of Species*, the scientific community had already become well aware of the existence of fossils that bore little resemblance to the extant animal kingdom. This gave rise to a debate over whether we ought to explain the natures of these past creatures by the present physical, chemical, and biological laws. Some disputants, fearing the oncoming paradigm shift in our understanding of world history, argued that it is illegitimate to apply present laws to explain past facts. *For all we know*, they argued, *nature is irregular, and so such extrapolation cannot be justified.* In fact, one of the crucial steps to accepting the Darwinian paradigm shift was accepting that present physical, chemical, and biological laws are applicable to the past. This was argued on the grounds of simplicity.

happened in the past, then we must be assuming what we had hoped to prove.

I am hoping that this rough presentation is enough to give the sense that there is some problem lurking under the attempts to justify induction. At any rate, it is enough for us to attempt to give an explicit reconstruction of the skeptical reasoning. To this end, I will follow David Hume (as closely as possible[4]).

Here is how I reconstruct Hume's skeptical argument:

1 All inductive arguments presuppose the principle of the uniformity of nature.

(Remember, this principle says that unobserved causes-and-effects resemble observed causes-and-effects.)

2 In order to justify the conclusion of an inductive argument, one must justify—that is, argue for—the principle of the uniformity of nature.

3 There are only two ways to justify the principle of the uniformity of nature: (i) using a "demonstrative" proof (a deductive argument with a priori premises) or (ii) using a "probabilistic" argument (an inductive argument with a posteriori premises).

4 You cannot justify the principle according to (i).

This is because, according to Hume, if the uniformity of nature could be demonstrated a priori, then it would not be possible for it to be false. But, he says, if it were *impossible* for nature to be non-uniform, then there must be some contradiction that arises if we were to suppose it to be non-uniform. However there's no contradiction in supposing that nature is non-uniform. He thus concludes that the principle cannot be demonstrated.

5 You cannot justify the principle according to (ii), because that would require reasoning in a circle.

[4] There are considerable questions about how to interpret Hume, which I will largely ignore. It is not my intention to give the most faithful reconstruction of Hume's original text; it is more important to me to outline the argument as it survives in the popular imagination. For instance, there is a substantive question as to whether Hume originally meant for his skeptical conclusion to imply that humans are positively *irrational* for relying on inductive arguments. Sometimes he seems to suggest so; other times, he seems to be making a descriptive point about human psychology, instead of a normative point that condemns our beliefs. According to the psychological reading, Hume is claiming that our beliefs about the future, and about cause and effect, arise out of our *habit* (our animal nature), rather than some divine faculty of reason. But this is merely a claim about how our minds are constituted; it is not to pass judgment upon these beliefs. I myself tend to read Hume this way. He was, above all, interested in human psychology (at a time before psychology had its own methods distinct from philosophy). But for this chapter, I will pretend that Hume is interested in defending the negative normative conclusion that our beliefs based on induction are *irrational* (in addition to being habitual).

This, as we said, follows from (1): that *all* inductive arguments presuppose the principle of the uniformity of nature. So if we were to try to employ an inductive argument to justify the principle, we must thereby presuppose what we're seeking to conclude.

6 The principle of the uniformity of nature cannot be justified (from premises (1-5).

C Inductive arguments cannot be justified (from 1-6).

2 Some Preliminary Problems with the Hume-Russell Problem of Induction

We have now uncovered our infamous skeptical foe. It cannot be emphasized enough how disastrous this conclusion would be if it were true. As I remarked in the opening, nearly all of our scientific beliefs come, in one way or another, from inductive methods and reasoning. This conclusion says that these beliefs aren't even *justified*. Nevermind whether they count as knowledge; this skeptic says that these beliefs are no better than guesses, wishful thinking, or deranged illusions!

Fortunately for us, there are many things fishy with the above argument. I have many minor quibbles, but I will leave them to a footnote.[5] For this section, I would like to state three major questions and criticisms of the Hume-Russell version of the problem of induction. I will not, at this time, claim that any one of these problems refutes 'the' problem of induction, because, as I have already cautioned the reader, I do not believe that there is only *one* problem of induction. Instead, I take these criticisms to point us in the direction of where the argument needs room for improvement. They will show what the skeptic must address if they are to marshall a credible case against inductive reasoning.

[5] (1) Hume's argument against the possibility of a demonstrative proof is unsound. It relies on the premise that *if P is impossible, then P implies a contradiction.* But there are impossible propositions that do not entail contradictions. Case in point: it is impossible for water to be anything other than H_20. But as a matter of sheer logic, there is no contradiction entailed from the proposition that water is not H_20. However, I say that this is a minor criticism because, regardless of the flaw in Hume's sub-argument, the premise still seems true enough. It still seems true that we cannot demonstrate a priori that nature is uniform. (2) Hume says that there are *only* two ways to justify the principle of the uniformity of nature. But why only these two? Why can't there be a deductive argument based on a posteriori premises? Or, why can't there be an inductive argument based on a priori premises? I'm not sure how this would go, but Hume hasn't given us any reason to think that his dilemma is exhaustive.

In one of the most famous episodes in the history of philosophy, Immanuel Kant claimed to have found a third way to justify the principle of universal cause-and-effect. He claims that the principle is both a priori *and* 'synthetic'—i.e. that its negation is non-contradictory. However, his defence relies on the crazy idea that the principle of cause-and-effect is not a description of how worldly objects behave "in themselves", but instead a matter of how *we* conceptualize and experience them. It is thus an artifact of our conceptual scheme that we *project* onto the objects we experience. I will not entertain this radical—and frankly, desperate—solution here.

2.1 Presuppositions and necessary conditions for justification

The first premise states that the principle of the uniformity of nature is a presupposition of all inductive arguments. Remember, according to the way that I've defined the term, a presupposition means that one must actually be justified in believing the principle in order to be justified to use induction. It isn't enough for the principle to be true—for nature to actually be uniform. One must also believe it with justification.

But why should we accept such a strong premise? Time and time again I have argued that we mustn't confuse the necessary conditions for knowledge or justification with full-blown presuppositions. That is, even if we grant that nature must be uniform for our inductions to be justified, it is still a wholly different matter to claim that we must justify this principle as well. Even if the *truth* of the principle is needed for justified induction, it still doesn't follow that *justified belief* in the principle is needed. So again, why does the skeptic think that induction has this stronger requirement?

We can put the matter another way. Let's grant, for the sake of argument, that inductive inferences are justified only if they are reliable, and that they are reliable only if nature is uniform (that is, uniform in certain respects that the principle of the uniformity of nature would specify). Imagine a person, Charles, who throughout his life has seen thousands of cats and has observed them to be middling size. Charles extrapolates that the next cat he sees will be middling size. Moreover, since the cat species has genetic parameters that determine each cat to be middling size, Charles' inference is highly reliable. It is extremely probable that the next cat that he sees will be middling size. Now contrast this with another subject, James, who has also seen thousands of cats. But unlike Charles, James reasons that *because he's seen so many mid-sized cats, he's due to see a cat that's very tiny or very large.* He arbitrarily decides that the next cat that he sees will be over a thousand kilograms. In other words, James reasons 'counter-inductively'; he expects that the patterns that he has observed will flip in the next instant. Of course, James is dead wrong about this. His inductive argument is highly unreliable, and it's exceedingly unlikely that his conclusion will be right.

Finally, let's imagine that Charles and James are ordinary people who haven't given too much thought into the philosophical foundations for their reasoning practices. In particular, neither of them has ever contemplated the principle of the uniformity of nature. Charles has never thought to defend it and James has never thought to deny it. So *a fortiori* Charles has not formed a *justified belief* in the principle, since he hasn't formed any belief about it at all.

Well, if Hume's first premise is right, then Charles' reliable induction is no more justified than James' crazy counter-induction. Charles is not believing any more rationally than James is, even though James' inference is clearly mad.

But how plausible is this, really? I would submit that it is exceedingly *less* plausible than Hume's first premise. It seems much more commonsensical to say that Charles is justified from having made a reliable inference than it is to say that Charles isn't justified because he hasn't thought about the principle of the uniformity of nature. (Recall the Moore-shift from §2.2.)

I won't pretend that this comment decisively refutes Hume's first premise since all that I'm doing is tugging at the reader's intuitions. But it does show that Hume needs more of an argument. He doesn't get his first premise for free, since it is, in fact, contrary to common sense. And since he's the one who's claiming to have revealed a devastating problem for induction, he's the one who ought to defend his premises.

2.2 A problem of deduction?

The second problem with the Hume-Russell problem of induction is that it doesn't explain why *induction*, in particular, is supposed to be especially problematic. It doesn't tell us what it is about induction that singles it out as an illegitimate form of reasoning.

To press this point, consider whether there is an analogous problem for deductive inference. Suppose that I were to essay a deductive argument: *if it rains tomorrow, then I ought to pack a raincoat; it will rain; therefore I ought to pack a raincoat.* And suppose now that a *deduction* skeptic sets out to challenge my inference. They say that if I'm to be justified in believing the conclusion, then I ought to justify my mode of inference as legitimate. If I cannot do this, then my inference does not render my conclusion justified.

Let's see what happens if I were to acquiesce to their challenge. It is well-known that deductive arguments of the *modus ponens* form (like the one above) are truth-preserving. This means that if the premises are true, then the conclusion must be true. So I could rightfully respond that *if my premises are true in this case ("if it rains tomorrow, then I ought to pack a raincoat", "it will rain tomorrow"), then so is the conclusion ("I ought to pack a raincoat"); the premises are true; therefore, the conclusion is true.*

No doubt the deduction skeptic will think that they've caught me in a trap. The argument I have just given is itself a deductive argument. In fact, it is an instance of *modus ponens*, the same kind of argument as my original. So if the point is to justify my original use of deduction, then my justification exhibits a kind of circularity.

Not only that, but the circularity involved in my justification for deduction is the same kind of circularity that was supposed to be so problematic in the induction case. In the original problem, it is claimed to be a mistake to justify induction, via the principle of the uniformity of nature, by *using induction.* But if that's a mistake, then it should also be illegitimate to justify *modus ponens* with a *modus ponens* argument, or justify deduction using deduction.

In both of these examples of deduction and induction, it might very well happen that the circular inferences are valid or reliable. My circular *modus ponens* argument is *in fact* valid (its premises *do* guarantee the truth of the conclusion), and it could turn out that the inductive justification for induction is reliable and delivers a true conclusion as well. But as we observed in the previous subsection, this isn't enough to satisfy the skeptic. The induction skeptic does not merely want a defence of induction that is reliable with a true conclusion; they want a defence that is *justified*. And reliability is not supposed to be enough. Well, in that case, why should the skeptic settle for less when it comes to deduction? Why should they accept a valid defence of deduction, if it involved an analogous circularity?

All of this is to make the point that induction-specific skepticism, in the style of Hume and Russell, appears to be on equal grounds with the analogous deduction-specific skepticism. Both kinds of skeptics can make the same moves and press the same points. If the Hume-Russell problem of induction were raising a genuine problem, then there must be a problem for deduction too. There's no reason for one to subscribe to one skeptical argument without subscribing to the other skeptical argument. Either both skeptical arguments work or neither of them do.

To be clear, I am not raising this point in order to shed any skeptical doubt upon deduction. Rather, I take it to be patently obvious that deductive inference can extend the frontiers of our justified beliefs. If I am justified to believe that P and *if P, then Q*, then I am thereby justified to believe that Q. Moreover, the reason that I am justified to reason in this way is precisely because it is a valid form of inference. Nevermind whether I can justify this claim to the satisfaction of the skeptic. (They are committed to never being satisfied and are hence being unreasonable). Deduction skepticism is clearly a false view.

This now puts the ball back into the induction skeptic's court. If their argument really is on equal footing with the skeptical argument against deduction, then that is a problem *for their argument*. Since the latter argument is flawed, their argument must have an analogous flaw. (The mistaken premise in the skeptical argument against deduction is the supposition that a deductive inference only confers justification if the subject can non-circularly argue that deduction is valid. By analogy, this sheds doubt on the second premise of Hume's argument.) So if the induction skeptic is to avoid this negative verdict, they must amend their argument so that it doesn't collapse into an implausible one-size-fits-all regress argument. Specifically, they must explain how induction differs from deduction in its justificatory features.

2.3 The principle of the uniformity of nature is too general

The final—and to my mind, most fatal—flaw in the Hume-Russell problem of induction has to do with the principle of the uniformity of nature itself. This principle is too vague and too general for the premises to be anywhere near plausible.

The fact that there is so much variation in how the principle is defined should already raise a red flag. As we have seen, Hume originally defined it as the thesis that unseen causes-and-effects resemble observed causes-and-effects. Other times the principle is defined as saying that the universe is governed by fundamental laws of nature. On other occasions, it is defined as the view that nature, as a whole, is regular or uniform, as opposed to random. Lastly, Russell offered it as the simple claim that the future will resemble the past.

Which is it? The various versions of the principle are not equivalent. And yet, the first premise of Hume's argument states that *the* principle must be justified in order for induction to be legitimate. This raises the question of *what* exactly is this principle that is allegedly presupposed by all inductive arguments. Hand-waving over this issue is a serious problem for induction skepticism.

However, this isn't even the biggest problem. Perhaps this worry can be settled by taking a single definition of the principle and sticking to it. Let's say (albeit, vaguely) that according to the principle, the entirety of nature follows regular and uniform laws.[6] Well, in that case, I would like to contend that this is far too general for the premises of Hume's argument to be true.

Under this proposed definition, the first two premises of Hume's argument jointly state that:

> For any inductive inference, if I am to be justified in believing its conclusion, I must first justify the claim that *all* of nature follows regular and uniform laws.

But do I really need to presuppose that *all* of nature follows uniform laws in order to have a justified inductive inference about some specific topic?

Take, for example, my inference that all cats are middling-size. To be licensed to draw this conclusion from particular observations of cats, do I really need to presuppose that, say, all *dogs* are uniform? Must I presuppose that the penguins are uniform? Does my inference depend on the uniformity of asteroids? Or the taste of tea? Or the price of bitcoin? Why on earth would this inference presuppose the uniformity of anything other than cats?[7]

Indeed, it seems as plain as anything that this inference will be justified even *if* other remote parts of nature are highly irregular or random. Imagine that there's a species of deep sea fish that has totally random size patterns. Some members of its kind grow to be as small as a dime while others grow to be as large as a whale, and there's no rhyme or reason to how they grow. Regardless, would this possibility in any way detract from my inference about cats? Must I rule it out in order to infer that *cats* are all middling-size?

It strikes me as absurd to think that particular inductive arguments, about a specific kind of thing, presuppose a *general, across-the-board uniformity* of *all* things. A cogent inductive argument about a kind of thing K at most only presupposes that the Ks are uniform.

Moreover, a cogent inductive argument *only* assumes uniformity in the aspects that are relevant to the property that is being inferred. When I infer that cats are middling-size, I only need to assume that cats are uniform with respect to their *size*. I do not presuppose that they are uniform with respect to their personality, or their fur colour, or the day of the week on which they were born. Only some kinds of uniformity are relevant, and they are quite specific to the subject matter of the particular inductive inference that I'm employing at the

[6] Hume's own formulation in terms of causes and effects is insufficiently general because not all inductive arguments are directly concerned with causation.

[7] And the laws of physics and biology that keep sizes stable.

moment. Again, an inductive argument does not, and should not, presuppose across-the-board uniformity. Only certain, highly-specific kinds of uniformity are relevant to a given inductive inference.

All of this goes to show that the first two premises of Hume's argument are false as they stand. We do not need to justify a general principle of the uniformity of nature. If there are any presuppositions for inductive inference, they must be something else. If the skeptic still has a case, then they must revise their argument accordingly.

When we first motivated the skeptic's concern with induction, we observed that each inductive inference makes certain assumptions about how the world works. (An inductive argument about the size of cats assumes that its sample is representative of all cats, and that the species of cats will be relatively uniform in size.) So, this much is true: every inductive inference makes an assumption of uniformity. But then the skeptic goes on to assert that there is an assumption of uniformity that every inductive inference makes. This is how they convert their observations about each inductive argument into a skeptical argument that applies to all of them. But it is also an unjustified leap in logic. From the fact that each induction makes an assumption of uniformity we cannot conclude that there is a single assumption of uniformity that every induction makes.[8]

I fear that the induction skeptic may have committed this logical fallacy. This error may explain how they came to their first premise. For, as we have seen, it's not as if their first premise is plausible on its own merits.

3 The New Riddle of Induction

We have now raised three worries about the Hume-Russell problem of induction as it was explicitly reconstructed. But the problem of induction is a slippery foe. Whenever it is refuted in one of its explicit forms, it tends to morph into something else.

The first criticism is that the Hume-Russell problem doesn't defend its claim that the subject must actually *show* that induction is reliable. Why isn't it enough for their induction just to *be* reliable? Secondly, it doesn't explain what is special about induction (as opposed to deduction) that makes it distinctly problematic. Thirdly, it is implausible to claim that there is a single principle of the uniformity of nature that underwrites every cogent inductive argument.

Bearing these criticisms in mind, it is possible to develop another version of the problem of induction that evades each of these difficulties. This is called the "New Riddle of Induction" and it comes from Nelson Goodman. In this section, I will explain the new riddle and explain how it is distinct from the old one. Doing so will give us a deeper understanding of the motivation behind induction skepticism, and point us in the direction of how to avoid it.

The best way to develop the new riddle is to first focus on an example of an inductive inference that we would ordinarily take to be good and justified.

[8]Compare this with another invalid argument: each country has a capital city, therefore there is a city that is the capital of every country.

To borrow Goodman's own example, let's say that we have inferred that all emeralds are green from a vast survey of green emeralds.

1 Emerald # 1 is green.

2 Emerald # 2 is green.

...

N Emerald # N is green.

N+1 No observed emeralds are non-green.

C All emeralds are green.

Imagine that N is a very large number; it can be as large as you'd like for a survey to be comprehensive. Now, if *any* inductive argument is legitimate and justifies its conclusion, then *surely* this one does. If it can be argued that this argument (in particular) is somehow bad, then induction is in trouble.

Goodman's new problem aims to show that even this argument could fail to justify its conclusion. However, he doesn't argue this by proposing any general, sweeping principles, like Hume and Russell did before him. Instead, he argues this by parody.

In the above argument we inferred that all emeralds are *green.* But instead of inferring that they're green, we could have used a different property. Goodman famously introduced the 'property' of being *grue.* It is defined like so:

> An object is *grue* if and only if it is either (i) green, and observed before 2100, or (ii) blue, and not observed before 2100.

Now, admittedly, grue is a fairly odd property.[9] But nothing stops us from defining it as we did. It is also entirely possible that our conceptual scheme contains concepts for properties that are similar to grue, and that upon further analysis, we find that they are mismatched combinations of dissimilar properties. For this reason, we need to take the idea of grue seriously.

But now there's a catch. Every emerald that we have observed has been green. So it follows from the definition of 'grue' that every emerald that we have observed is also grue! Whenever we have observed a green emerald we have thereby observed a grue emerald. Therefore, we can amass a list of true premises that correspond to the premises of our original argument. We are thus confronted with another potential inductive inference:

$1'$ Emerald # 1 is grue.

$2'$ Emerald # 2 is grue.

[9] I put the word 'property' is scare quotes because one might object that this isn't *really* a property: it is a grotesque amalgamation of properties. Fair enough; there may be something to this objection. But let's bracket that worry for now. Think of properties as being fairly permissive. Anything counts as a property as long as you can define it.

...

N′ Emerald # N is grue.

N+1′ No observed emeralds are non-grue.

C′ All emeralds are grue.

Again, each one of these premises is *true*. Moreover, they are exactly the same in number as the premises of our original argument. So, besides the use of 'green' in one argument and 'grue' in the other, the two arguments are otherwise the same. We might put this point by observing that they both have the same logical *form* or *structure*. That is, they are each instances of this common archetype:

1* Instance # 1 of K is P.

2* Instance # 2 of K is P.

...

N* Instance # N of K is P.

N+1* No observed Ks are non-P.

C* All Ks are P.

But despite all of this similarity, their conclusions conflict with one another. The original argument implies that an emerald observed in 2100 will be green, whereas this parody argument implies that an emerald observed in 2100 will be blue.

This now presents a problem for the advocates of our original inductive inference. The trouble is that the parody argument is patently *bad*. Clearly, we *shouldn't* infer that all emeralds are grue. We should not expect to see blue emeralds after the year 2100. But if the inference to all emeralds being grue is unjustified, then perhaps the original argument is unjustified as well.

We can even put this point into an explicit skeptical argument debunking any inductive inference that we please.

1 Suppose that we make an inductive inference from N observations of objects of the kind K being P to the conclusion that all Ks are P.

2 Then we can define a grue-like property P* and construct a parody argument with all true premises *and exactly the same logical form* (same number of premises).

3 The two arguments will therefore be on par in all of the respects that are relevant to justification.

4 Thus, if the original argument confers justification, then so will the parody argument.

5 But the parody argument does not confer justification.

C Therefore, our original inductive inference will not justify the conclusion that all Ks are P.

Since this conclusion is perfectly general, it is sufficient to cast doubt on all inductive inferences. We thus have another sweeping problem of induction on our hands.

Before we turn to assessing this new problem, notice that it completely avoids two of the criticisms of the Hume-Russell version. Most saliently, this argument never mentions any alleged presupposition of a 'principle of the uniformity of nature.' It doesn't demand for the ordinary subject to argue, justify, or prove any abstract principle for their inductions to be justified. It thus completely avoids the first and third objection that we raised against the Hume-Russell problem of induction. As for the second objection, we will see shortly how this new version of the problem distinguishes between induction from deduction.

4 The Material Theory of Induction

There are really two ways that we can draw lessons from Goodman's riddle. On the one hand, we can see it as exacerbating the old problem of induction. This is how the riddle has been presented so far. On the other hand, we might also see it as providing a valuable insight into the nature of induction and the vast difference between induction and deduction. The latter route will open us up to seeing induction in the right way, which will pave the way forward to dissolving the skeptical worries directed towards it.

When we look at the explicit reconstruction of the Goodman-style problem of induction, the third premise ought to draw our attention. It says that the two arguments—the original and the parody—are on par with respect to justifying their conclusions. So if the original induction to "all emeralds are green" is justified, then so is the parody induction to "all emeralds are grue", and vice versa.

We must take a closer look as to why this premise might appear to be true. The answer, I take it, is that the original argument and the parody argument have this in common: they each have justified (and true) premises and they exhibit the same logical form. Remember, they have the same 'logical form' because they each have the same number of premises and their premises and conclusion have the same grammatical structure. The only difference is that they invoke different properties, *green* and *grue*. It is for this reason that the Goodman-style induction skeptic takes the original argument and the parody argument to stand or fall together.

This brings to light an assumption that is driving the skeptical conclusion of the new riddle of induction. It assumes that the cogency of an inductive inference depends solely on two features: (i) the justification of the premises and (ii) the logical form of the argument. It is this assumption that underlies the claim that the "green" induction and the "grue" induction share the same cogency. If this assumption were false, then the cogency of each respective argument may depend on other factors. Perhaps, in that event, we could say that there's a relevant difference between the property of *being green* and the property of *being grue* that makes the original justified while the parody augment unjustified. If so, then we could escape the skeptical conclusion of Goodman's riddle.

So why would the skeptic be tempted by this assumption? Why would the justification of an inductive inference be *solely* a matter of logical form and the status of the premises? Here is my suspicion. It seems to me that the skeptic accepts this assumption because they are thinking of inferential justification on the model of *deductive* inference.

When it comes to deductive inference, the analogous assumption is actually true. The hallmark of deduction is that, for any valid deductive argument, *if all of the premises are true, then the conclusion must be true.* It is thus highly plausible to think that justification for the premises, combined with the valid logical form of the argument, is sufficient to grant justification for the conclusion. Moreover, the *validity* of a deductive argument depends solely on its form. For example,

1 All geese are swimmers

2 Albert is a goose

C Therefore, Albert is a swimmer

is a valid deductive argument. Since this argument is deductive and valid, any other argument with the same form will also be valid. This means that we can replace any category term with "geese" or "swimmers", or any individual term with "Albert", and yield another valid deductive argument. Say we replace "geese" with "moon rocks", "swimmers" with "grue", and "Albert" with "Carl":

1′ All moon rocks are grue.

2′ Carl is a moon rock.

C′ Therefore, Carl is grue.

We thus obtain another valid argument. And if, *somehow*, we could justify the premises, we would thereby justify the conclusion. This is somewhat fanciful because I have no idea how we could justify the claim that *all moon rocks are grue*, but regardless, if we could somehow do it, we would thereby justify the conclusion that Carl the moon rock is grue.

All of this goes to illustrate that for deductive arguments, their justificatory status depends entirely on their premises and their logical form. You can, so to speak, read off the epistemic features of a deductive argument just by looking at its structure and the justification for each premise, taken individually.

Incidentally, this feature also explains why a deductive inference cannot be debunked by a Goodman-style parody argument. If you have a good deductive argument with justified premises, and you replace them with other *justified* premises that invoke grue-like properties while leaving the form intact, then the resulting 'parody' argument will *still* be a good argument! The upshot of this is that, unlike the Hume-Russell version of the problem of induction, Goodman's riddle has identified a puzzle that is unique to induction. There is no analogy of the new problem of induction for deductive inference.

This feature of deduction is certainly a very nice thing. It means that deductive inferences are extremely well-behaved and they can be straightforwardly checked for their epistemic properties. It would therefore be desirable if all inferences could work so mechanically. But Goodman's riddle dashes our hopes for this. His argument shows that we cannot have both (i) justified inductive inferences and (ii) a *formal* procedure for checking their justification. Now, it seems to me that the skeptic is motivated by this outcome to cast their doubt upon (i). With the contrast to deduction in mind, we may reconstruct their motivation like this. First, they see that the justification for deduction depends entirely on the justification of the premises and the logical form of the argument. Then they figure that inductive inferences ought to receive their justification in the same way. Hence, they infer that *if* any inductive inferences are justified, then they ought to be justified solely in virtue of the premises and the logical form. But this isn't the case for induction; Goodman's parodies show that induction doesn't work like that. The skeptic then concludes inductive arguments—unlike deductive arguments—never justify their conclusions.

This way of casting the skeptic's motivations should make it abundantly clear how we ought to respond. The key is to stress that *induction is not at all like deduction*. Unlike deductive inference, an inductive inference, when it is justified, is not solely justified by its premises and its form. To be specific, we ought to reject the third premise of the new problem along with the hidden assumption that underlies it.

This, of course, raises the question: if cogent inductive arguments aren't justified solely in virtue of their premises and their form, then what else does their justification depend on? To answer this, it is important to look carefully at each individual case.

Consider again the contrast between the inference to "all emeralds are green" and the inference to "all emeralds are grue." Evidently, if there is to be any difference in the justification between these two arguments, it must depend on the properties that these arguments invoke. Specifically, it must depend on the nature of *emeralds* (as a general kind of stone) and the properties of *being green* and *being grue*.[10]

It was remarked earlier in this chapter that inductive arguments tend to rely on hidden assumptions about their subject matter. In particular, they depend on their samples being representative and on their subject matter being uniform in the relevant respects. Well, now we can make this point explicit in the green and grue cases. Each of these arguments must depend on hidden assumptions of uniformity in their subject matter.

[10]This again contrasts with the case of deduction. Once the premises of a deductive argument are justified, then the license to infer from the premises to the conclusion no longer depends on specific features of the subject matter.

However, the assumptions will be different for each argument. For the first argument, we rely on the assumption that *emeralds are likely to exhibit similar colour patterns in virtue of belonging to a common kind of gemstone.* Let's call this assumption the 'material postulate' since it is a postulate about the *material* (i.e. the subject matter) of the inductive inference.[11] In this case, the material postulate is, in fact, *true.* It is a fact that gemstones tend to exhibit similar colour patterns within a common kind. That is because a kind of gemstone is grouped together by common chemical properties and therefore common light-reflection properties. The truth of this material postulate explains why the premises of the induction confer a high degree of probability towards the truth of the conclusion.

Contrast this now with the parody argument involving *grue.* What would the material postulate be for this argument? It cannot be the same assumption, because *grue* is not a colour. It is an amalgamation of colours indexed to our observations over time. So perhaps the material postulate for the grue induction is that *emeralds are likely to exhibit similar patterns of one-colour-at-one-time and another-colour-at-another-time, depending on when they're observed.* But such a material postulate is as false as anything can be. Emeralds are naturally occurring substances that keep the same chemical composition over time. They don't change their colour depending on the year or when they're observed. So this *false* material postulate does not underwrite a cogent inductive argument.

Notice that it is not only the predicate that figures into the material postulates of an induction. In our examples, the material postulates are not only concerned with *being green* and *being grue*; they are also concerned with the *kind* that is the subject of the induction. The kind must exhibit an objective *unity* or *similarity* amongst its objects in order to be fit for a cogent induction. To see what I mean by this, contrast the emeralds case with another superficially similar argument. Instead of emeralds, suppose we infer that *all jade stones are green.*

1 Jade stone # 1 is green.

2 Jade stone # 2 is green.

...

N Jade stone # N is green.

N+1 No observed jade stones are non-green.

C All jade stones are green.

Because of its formal similarity with induction to "all emeralds are green", one would think that this argument would be just as cogent. But in this case, there's a problem. What we call "jade" turns out *not* to be a single kind of mineral. There are actually two kinds of minerals—jadeite and nephrite—that

[11] This terminology comes from John Norton's "A Material Theory of Induction" which will be discussed shortly.

are chemically unrelated but happen to both be classified as jade because of their superficial similarity. So if we were the ones running this induction, and it turned out that all of our samples of jade were nephrite, then we would have no grounds to infer that the other kind of jade, namely jadeite, is green. Our sample wouldn't be representative because jade isn't naturally unified and we only encountered one kind of it.

All of this goes to illustrate the point that I've been leading on. Each inductive inference relies on an assumption that its subject matter is *unified* in the relevant respects. This is the material postulate of an induction. For an inductive argument to be *cogent*, its material postulate must (at bare minimum) be *true*. The subject matter must *really be* uniform in the respect that is suitable for the inductive inference. This is what separates the original 'green' case from the 'grue' case. The induction to "all emeralds are green" has a true material postulate: emeralds *really are similar* with respect to their colour properties. However the induction to "all emeralds are grue" has a false material postulate: emeralds are *not* similar with respect to the strange kind of properties that is represented by 'grue.' This explains why the former argument is cogent whereas the latter one is not. A similar thing can be said about the 'emerald' case and the 'jade' case. Emeralds *really* are unified into a common kind whereas jade stones are not. Hence there are cases where the former will admit of cogent inductive generalization whereas the latter won't: namely, when we only have a sample of one kind of jade.

It is important to notice that the 'unity' asserted by the material postulates of these inductive inferences is a thoroughly objective matter. By this, I mean that the class of emeralds is unified together by an objective chemical similarity. The unity is not merely a matter of our conceptualizing the emeralds under a common category. It is not *up to us* that the emeralds belong together. The world is carved up, independent of our concepts, into the stones that are emeralds and everything else. This makes it a purely factual matter, independent of our theorizing, as to whether the material postulates of our inductive inferences are true.

The idea of positing *objective unities* and *natural kinds* to solve Goodman's riddle has been a popular one ever since Goodman devised the problem. David Lewis is one example of a well-known proponent of this solution. The basic idea is that some properties and kinds are eligible for induction whereas other properties and kinds are not. *Greenness*, in virtue of the objective similarity among green things, is fit for cogent induction; *grue*, on the other hand, is an unnatural property and so it is no good. On the same token, emeralds form a natural kind, and are hence objectively good for inductive generalizations; whereas unnatural groupings, like the class of grue things, are bad, and only make for illicit inductions. Moreover, which kinds and which properties are eligible for induction is an objective feature of the world. It isn't up to us. Sometimes nature blesses us with justification and sometimes it doesn't.

I am broadly sympathetic to this idea, but only if we aren't misled by taking it too literally. I don't want to suggest that there is a single sense of 'unity' or 'natural kind' that can serve a one-size-fits-all analysis of the justification of induction. The Lewisian solution invites the picture that there are two kinds of properties—the natural ones and the non-natural ones—and hence we can analyze inductive cogency as *generalizing a natural property to a natural kind.* I, on the hand, do not wish to pretend that there's a single sense of 'unity' or 'naturalness' that underwrites inductive cogency. On the contrary, I believe that the material postulate of each inductive inference will vary from case to case.

The main precedent for the view that I'm advocating is the "material theory of induction" by John Norton. The main slogan for Norton's theory is that "all inductive inference is local." This means that what makes an inductive inference good, when it is good, is specific to the topic. Each inductive argument has its own material postulates, but the material postulates are specific to the subject matter of the induction. They each posit *some* kind of unity, but the respect of unity that's relevant will depend on what the argument is about and what it says about its subject matter. For a thoroughgoing material theory of induction, there really is no general formula for what makes an inductive argument justified (if it is justified). As Norton says, "the admissibility of an induction is ultimately traced back to a matter of fact, not to a universal schema."[12] But that is not a *problem* for induction; it's a feature that *distinguishes it from deduction.* The fact that there is no universal test for inductive cogency does not prevent every inductive argument from being cogent. Indeed, this theory allows for many inductive inferences to work, but each for their own reasons.

Let me give a couple more examples to illustrate the variety. Earlier we considered an inductive inference to the conclusion that all cats are middling-size. For this argument, the material postulate is that the category of *cats* form a natural biological class, and as such, they are likely to be similar to each other with respect to size. This is made true by the fact that the cats do belong to a common species, and therefore have a high degree of genetic overlap. Moreover, the genetic code for a species codes for the size of its organisms. Since all of this is true, it follows that my inductive inference is cogent; its conclusion is justified.

For another example, suppose that I take a handful of samples of water boiling at sea level, and observe that they always boil at a 100 degrees celsius. I thereby infer that all water boils at a 100 degrees celsius. This is another cogent inductive argument and I am justified in drawing this conclusion. This time, the material postulate says that *water is a common chemical kind, with an underlying chemical unity.* The similarity of each water sample consists of the fact that samples of water have the same chemical composition. Moreover, it is also a fact that the samples of any chemical kind are likely to share a common boiling point because boiling points are determined by the facts of chemistry.

I bring up these examples to drive home the point that the material postulates—the facts that ground the cogency of each induction—really are a heterogeneous lot. For the cat example, the unity of the kind (cats) and the stability of the

[12] Norton, J. D. (2003). A material theory of induction. *Philosophy of Science*, 70(4), 647-670.

property (size) advert to the facts of biology. For the water example, they advert to the facts of chemistry. Again, we shouldn't assume that there's any one kind of unity that grounds every induction. As we saw with the Hume-Russell problem of induction, that assumption is a mistake.

All of this talk of material postulates will no doubt arouse the skeptic and provoke them into attempting one final ploy. They might grant what I've said so far about the heterogeneity of the facts that ground inductive cogency. "Fine", they will say, "have it your way. There is no *one* assumption that is common to every inductive inference. But each induction has its own assumptions and *now you have to justify those.* How do you justify the material postulates? For if you don't justify those, then your inductive inferences aren't justified. And if you can never justify any of them, then no induction is justified! When you infer that *all Ks are P*, how do you know that the kind K is unified in the relevant way (that it is more like emeralds than jade)? How do you know that the property P is stable in the relevant way (that it is more like green than grue)?"

If any skeptic finds this objection tempting, then they haven't been paying sufficient attention to the theory of induction on offer. To request an answer, once and for all, as to how one justifies the material postulates of the various inductive inferences is to fail to appreciate that they come in all shapes and sizes. It is therefore unreasonable to request a general answer as to how they are justified. There is no universal answer to the question of *how do you justify the grounds for induction.* You have to take them case by case.

With that said, there are a couple comments to make about justifying the material postulates. First of all, there is no a priori reason to think that for every induction, its material postulate has the status of a presupposition—that is, a proposition that one must be justified in believing if the induction is to be cogent. They might instead operate as *necessary conditions* for cogent induction. In that case, the subject does not have to defend them; they just have to be *true.* This is akin to the externalist answer to the problem of induction that we entertained in §4.2.1. The idea is that, for a special class of inductive inferences, the subject does not need to *justify* the reliability of the induction; it just has to *be* reliable and then it will confer justification. I will not propose here any definite thesis as to which inductions have this status. But if I were to hazard an opinion, I would say it is those that involve concepts that are fairly central to our conceptual scheme as human beings: e.g. 'person', 'animal', 'food', 'water', etc.

Secondly, I will also allow that there are plenty of inductive arguments that are such that their material postulates are genuine presuppositions. These are inductions where one must justify a belief in the material postulates if they are to be justified at all. The material postulates don't just come for free. Even if the inference is reliable, that wouldn't be sufficient for justification.

Which inductive inferences are those? The first that come to my mind are inductive inferences that involve a kind of thing that we wouldn't have known had existed if it weren't for a lot of scientific (and inductive) inference. So inductions that involve electrons, quarks, and genes are a fairly obvious representative of this kind.

Is there any special skeptical worry about inductions whose material postulates require justification? No. For we may justify their material postulates by using other inductive arguments. And contrary to the Humean skeptic, there is no circularity here according to the material theory. It is not circular to justify one induction with another induction if the two inductions rely on different assumptions. Let's say that there's an inductive argument I_1 that is based on assumption A_1, and A_1 requires justification. Well, perhaps we can justify A_1 by induction I_2 that is based on assumption A_2. And perhaps we can justify A_2 by induction I_3 that is based on A_3, and so on. None of this would constitute circular reasoning. At no step in this process are we assuming (or otherwise relying on) what needs to be proved. Sure, each argument has the form of an inductive generalization, but according to the material theory, that is a superficial similarity. It isn't their *form, as inductions*, that makes them justified or not; it is their respective assumptions about their subject matter.

Is there any worry that this process will result in an infinite regress? Again, no. Because as I've said, there's no reason to suppose that the cogency of every inductive inference depends on its material postulate being justified by the subject. Some inductive arguments are good simply because their material postulates are true and the inference is reliable. In that case, there's no need for further inferences. The regress can rest at a stopping point.

5 Dissolving the problem of induction

In this chapter, I set out to argue that the problem of induction, despite all of its influence, has been largely overrated. Now that I have given my answer, notice that I have not tried to *solve* the problem of induction by offering a traditional answer. I have not granted to the skeptic that there is a single problem that can be met head-on. I do not respond to the skeptic by handing them a single formula that purports to sort out the good inductions from the bad.

Instead, I prefer to say that the problem of induction ought to be *dissolved*. We do this by pointing out that the skeptical worries that generate it rest on mistaken assumptions.

Both the Hume-Russell problem and the Goodman-style problem take it for granted that if any inductions are justified, then there must be some common feature to all of the legitimate ones that mark them out as justified. (For Hume, this is the principle of the uniformity of nature; for the Goodman-style skeptic, this is their logical form.) To respond to these skeptics, we reject the idea that this is how inductive cogency works. When the skeptic demands to see the universal rules that make induction arguments good, I respond: there are no *universal rules*, but that doesn't imply that *specific* inductions can't be good!

If I am right, then there is no sweeping problem of induction. But where does that leave us? It leaves us in the totally banal situation where some inferences are good, some are obviously bad, some are unwittingly bad, and some are tough nuts to crack. It leaves us with plenty of situations where there's no principled way to know which inductive inferences to draw and which ones not to draw. But as I said at the outset, these are problems for the philosophy of science, not for foundational epistemology. For the epistemologist, the main task is to ferret out the sources of misplaced skepticism. So given that these particular skeptical assumptions have been exposed and debunked, I believe that my job is done.