Chapter 34

Inductive Reasoning is Not Justificatory Reasoning (I)

1. Our key thought

Pick up a small object from your desk. (A pen, for instance, or an eraser. But not something that might break or spill.) Hold that object about 20 centimetres above the surface of the desk. Now release your grip on it. Take the object again in your hand, hold it about 20 centimetres above the surface of the desk, and again release your grip on it. Do this three more times. Now describe what you witnessed, beginning with the first time you took the object into your hand.

Your description will go something like this: five times I held an object above my desk and then released my grip on it; each time I released my grip on it, the object fell to the desk.

Now take that object in your hand and hold it over your desk once more. What do you expect will happen should you now release your grip? –You expect that the object will again fall to the desk. And you are extremely confident that your expectation of it falling will be borne out by the course of events should you release your grip on it.

Are you warranted in predicting the future as you have? What reason have you to think that the object will fall to the desk rather than, say, remain suspended in the air or fall in an arc to the floor or simply disappear? Is your confidence in your expectation reasonable? –You possess, you say, a good argument to the conclusion that the object will fall to the desk. And, you add, it is because you possess this good argument that you are perfectly reasonable in being confident that your prediction is correct.

Here is the argument you say backs up your prediction and makes your high degree of confidence in that prediction perfectly fine:

1. I have five times now released my grip on the object I held over my desk.
2. Each of the five times I released my grip on it, the object fell to my desk.
3. I am now holding the object over my desk for a sixth time.
4. Nothing in the situation around me has changed from the first five times I held the object and released my grip on it.

Therefore: 5. When I release my grip on the object for the sixth time, it will fall to my desk.

Now release your grip on the object for the sixth time and observe what happens. (The object falls to the desk, just as it did the previous five times.) Pick the object up once more.

Consider the following two arguments:

1. I have *six* times now released my grip on the object.
2. On each occasion, the object fell to the desk.
3. I'm about to release my grip on it for the seventh time.
4. Nothing in my situation has changed from the first six times I released my grip on it.

Therefore: 5a. When I release my grip on the object for the seventh time, the object will fall to the desk.

1. I have *six* times now released my grip on the object.
2. On each occasion, the object fell to the desk.
3. I'm about to release my grip on it for the seventh time.
4. Nothing in my situation has changed from the first six times I released my grip on it.

Therefore: 5b. When I release my grip on the object for the seventh time, the object will not fall to the desk (but will instead remain suspended in the air, or rise, or disappear, or anyways do something other than fall to the desk).

The argument with 5a as its conclusion is a strong inductive argument all the premises of which are true; the argument with 5b as its conclusion is what we will call a counter-inductive argument. To create a counter-inductive argument, simply take a strong inductive argument all the premises of which are true and negates its conclusion. (No one, of course, reasons counter-inductively. Counter-inductive arguments are fictions concocted by philosophers interested in exploring inductive reasoning.)

That you accept the premises of the first argument, you say, justifies you in accepting the conclusion of that argument. That you accept the premises of the first argument justifies you in confidently expecting that the object will fall when you release your grip on it. The second argument, the counter-inductive argument, though, has all and only the premises of the first argument. But those premises, you say, justify one in believing 5a, not 5b. The second argument is just crazy. It's conclusion is not supported by its premises.

But why, then, do we think that we have more reason to accept predictions based on strong inductive arguments than to accept the negations of those predictions? Are we right to think that we have more reason to expect the course of events to bear out the conclusions of strong inductive arguments with true premises than to expect they will bear out the conclusions of counter-inductive arguments? Would someone confident that 5b will turn out true, someone whose confidence was based on the second argument, be making a mistake? What mistake? Do we have a good reason for thinking that his confidence is unreasonable or misplaced?

Our key thought in answering these questions must be that strong inductive arguments with true premises can be relied upon to generate predictions that will be borne out by the course of events. Our confidence in these predictions is well placed because of this reliability. Corresponding counter-inductive arguments, we add, cannot be relied upon to generate predictions that will be borne out by the course of events.

2. Strong inductive arguments

Strong inductive arguments with true premises can be relied upon to generate predictions that will be borne out by the course of events, or so we think. But in virtue of what is an inductive argument inductive? And in virtue of what is a strong inductive argument strong? Let's make sure we know the answers to these questions before returning to examine our key thought.

Our concern is not with inductive arguments generally, but only with those inductive arguments the conclusions of which are about what will happen some time hence (or what would happen some time hence were something else to happen first). That is, the arguments

with which we are concerned generate predictions. Given a prediction we see to be generated by an inductive argument we judge to be strong and all of whose premises we accept as true, we tend to expect that that prediction will be borne out by the course of events. For our purposes, all we need note in trying to understand what makes an argument inductive is that an inductive argument that has a prediction as its conclusion is inductive in virtue of its conclusion not following deductively from its premises. Generally, that means that the premises will concern the past or the present, not the future, and none of them will be a universal statement assumed to cover future events as well as past and present ones.

So, the arguments with which we are concerned state in their premises that events of one type have been, so far as we know and to this point in time, positively correlated with events of another type. And they conclude that an event of that second type will occur should another event of the first type occur. Let's try to make this more concrete in an example.

Here is a strong inductive argument (don't worry about the fact that its premises are false):

1. Many hundreds of thousands of widgets with black parts have been exposed to water in many different sorts of condition. (Widgets have been exposed to water when it's been windy and when it's been still; in the presence only of children, in the presence only of adults; at high altitudes, at low altitudes; in forests, in pastures; from above, from below; with liquid water, with ice; some have been exposed to a little water, some to a lot of water.)
2. Every time a widget has been exposed to water, all of its black parts have quickly turned red.

Therefore: 3. The black parts on the next widget that is exposed to water will quickly turn red.

Here, by contrast, is a weak inductive argument (again, ignore the truth-value of its premises):

1. Just three widgets with black parts have been exposed to water, and these three widgets were together in a group when they were all soaked from above with water from the same pail, and that was years ago and it happened thousands of miles away from here.
2. The black parts on two of those widgets quickly turned red after exposure to water while the black parts on one widget remained black.

Therefore: 3. The black parts on this widget will turn red soon after I spray this widget lightly with water.

In both arguments the premises concern the past and the present and the conclusion is a prediction about what will happen in the future (or what would happen in the future were some condition fulfilled). Neither argument contains a generalization such as that the black parts on any widget turn red when exposed to water (from such a generalization the conclusion follows deductively). And so they are both inductive arguments.

But one of the two arguments is much better than the other, or so at least we might well think. The strong inductive argument above has at least three features lacking in the weak inductive argument. 1) The premises make reference to many instances in which events of the first sort occurred. The premises tell us that the black parts of widgets have many times

been exposed to water. 2) The premises tell us that events of the first sort occurred in all kinds of conditions. They tell us that the black parts of widgets have been exposed to water in very many different circumstances. 3) The premises tell us that an event of the second type so far has always followed the occurrence of an event of the first type. They tell us that always the black parts turned red after the widget was exposed to water. Never has it happened (if the premises are true) that a widget's black parts did not turn red after exposure to water.

It would seem to be the presence of these three features that accounts for our confidence that the strong argument's conclusion would be true were that argument's premises true, a confidence we lack in the case of the weak argument. If we believe the premises of the strong argument, we will expect that the black parts on the next widget exposed to water will turn red, and we will quite confidently expect this to happen. On the other hand, even if we believe the premises of the weak argument, we will not be inclined to accept its conclusion. Maybe the black parts will turn red, but maybe they won't. Given only these premises, we don't think we can say what will happen.

Generally, then, this is the form of a strong inductive argument that has a prediction as its conclusion:

1. Events of type A have very often occurred.
2. They have occurred in a great many different circumstances.
3. Whenever an event of type A has occurred, an event of type B has occurred.

Therefore: 4. When an event of type A next occurs, an event of type B will also occur.

Strength, of course, comes in degrees. Some strong inductive arguments are stronger than other strong inductive arguments. The strength of the argument increases as more events of the prior type are covered in the premises, as the range of circumstances in which they have occurred increases, and the higher the proportion of times events of the target type followed events of the prior type. And, psychologically, our confidence that the prediction stated in the conclusion of the argument will be borne out is greater the stronger we think the argument is (given we believe the premises to be true, of course).

Now back to our question. We are, as a matter of fact, confident that the future will turn out as predicted in the conclusion of an inductive argument we take to be strong and all of whose premises we accept as true. As a matter of fact, when we possess an argument we take to be strong and all of whose premises we accept, we expect that the future will go as described in that argument's conclusion. Are we, though, justified or warranted or reasonable either in expecting what we do or in expecting it as confidently as we do? We would be justified or warranted or reasonable were we justified or warranted or reasonable in believing that we can rely on the conclusions of strong inductive arguments the premises of which are true to be borne out by future events. But can we be justified in believing that the conclusions of strong inductive arguments will be borne out by future events?

3. Inductive reasoning is not justificatory reasoning

Here is an argument meant to go some distance toward establishing that inductive reasoning is not justificatory reasoning. Here, that is, is an argument meant to go some distance toward establishing that we can have no good reason to think that the conclusions of strong inductive

arguments will borne out by the course of events, even those strong inductive arguments that have no false premises. It is a fairly detailed argument, so go through it slowly and carefully.

1. In order to show that we have better reason to accept as true predictions generated by strong inductive arguments with true premises than we have to reject those predictions, we need to show that strong inductive arguments with true premises can be relied upon to generate predictions that will be borne out by the course of events.
2. In order to show that strong inductive arguments with true premises can be relied upon to generate predictions that will be borne out by the course of events, we need to find a good argument that has as its conclusion that strong inductive arguments with true premises can be relied upon to generate predictions that will be borne out by the course of events.
3. Any inductive argument to this conclusion ("strong inductive arguments with true premises can be relied upon to generate predictions that will be borne out by the course of events") will employ the very method of reasoning whose reliability is at issue, namely, induction.
4. To argue that a method of reasoning can be relied upon to generate true conclusions by using that very method of reasoning is to beg the question of that method's reliability.
5. A question-begging argument in favour of a claim provides no reason to think that that claim is true.

Therefore: 6. We cannot find a good inductive argument that has as its conclusion that strong inductive arguments with true premises can be relied upon to generate predictions that will be borne out by the course of events.

Therefore: 7. That a prediction is generated by a strong inductive argument with true premises gives us no more reason to accept that prediction as true than to reject it as false.

Statement 7 in fact does not follow from statement 6; can you say why it doesn't follow? We'll come back to this point and tidy up the argument later.

The argument, more informally, is just this. If we are to be justified in believing that predictions generated by good inductive arguments are likely or typically true, we will have to have a *reason* for believing that predictions generated by good inductive arguments are likely or typically true. But we cannot have a reason for believing this. And so, even though we might believe it, we are not justified or warranted or reasonable in believing it. To have a reason for believing induction is reliable, we need to have a good argument that has as its conclusion that induction is reliable. Any inductive argument to this conclusion, though, will beg the question by presuming that induction is reliable. And so we cannot show ourselves justified in believing that induction is reliable through arguing inductively that it is.

Should we accept the premises of this argument? Let us look more closely at statements 3, 4, and 5.

Statement 5 is that a question-begging argument is not a good argument. One cannot defend a claim on the basis of an argument that begs the question in favour of that claim. An argument begs the question when it somehow presupposes the truth of the claim at issue.

Here's an example. Suppose someone tells you that Ivan is trustworthy. You ask what reason we have to think that that is true, that Ivan indeed is trustworthy. You're told

that since we can rely on Ivan's word and Ivan has assured us that he is trustworthy, we can be confident that Ivan is trustworthy. –Wait a minute!, you exclaim. –Those premises give reason to believe that Ivan is trustworthy only if we already suppose that Ivan is trustworthy, for otherwise Ivan's assurance that he is is worthless. The premise that we can rely on Ivan's word smuggles into the argument the very claim that is at issue. Thus the argument simply begs the question of Ivan's trustworthiness. You rightly dismiss the argument as entirely unhelpful: it gives you no reason at all (not even a little reason) to think Ivan is trustworthy. (Perhaps we should note that none of this gives you reason to think that Ivan is *not* trustworthy. The question whether Ivan is trustworthy is entirely still up in the air.)

An argument that begs the question in favour of its conclusion gives one no reason at all to accept that conclusion as true. If we accept this, that question-begging arguments are not justificatory arguments, our next concern will be with whether an inductive argument in favour of induction must beg the question in favour of induction as a form of justificatory argument. Statements 3 and 4 together imply that inductive defences of induction are inevitably question-begging. Are they?

An inductive argument meant to show that the conclusion of a good inductive argument is likely true will proceed as follows. Consider that predictions generated by strong inductive arguments with true premises have so far most often been borne out by the course of events. But then, since to date such predictions have typically come true, we can conclude that strong inductive arguments with true premises will continue to generate predictions that will be borne out by the course of events.

This argument is an inductive argument, certainly; the form of inference that takes us from premise to conclusion is induction. And it is a strong inductive argument whose premise is true (let us agree). But, then, this argument gives us reason to think its conclusion true only if we already accept that good inductive arguments tend to have true conclusions. But whether good inductive arguments tend to have true conclusions is just what we are attempting to determine. And so our inductive argument that good inductive arguments tend to have true conclusions begs the question that they do.

An analogous argument might help us see more clearly the point of the preceding two paragraphs. The sort of inductive reasoning with which we are concerned is not the only method people use to generate predictions or develop expectations. Some people cast horoscopes, some listen to their hearts, some consult the chicken oracle, some go with celebrity opinion. Imagine someone who forms expectations by listening to his heart or by asking her favourite celebrities setting out to discover whether he or she is well justified in using that method. How will he or she investigate the issue? What the heart-listener wants to know is whether he has a good reason for conforming his expectations to the deliverances of his heart. He wants to know whether his heart informs him truly. So, employing his favourite method of fixing expectations, he asks his heart whether his heart will inform him truly. Suppose his heart says "yes." Has this person discovered a good reason to think that his heart will inform him truly? Not at all, we say. He has begged the question in favour of his method by using that very same method in investigating that method.

The person who reasons inductively to discover that induction is trustworthy is no different than the person who consults his heart to discover that heart-consultation is trusthworthy. If one's heart's telling one that heart-consulting is fine is no reason to think

that in fact heart-consulting is fine, then one's having an inductive argument that inductive reasoning is fine is no reason to think that in fact inductive reasoning is fine.

The two cases, one might note, differ from each other in this respect. The person who listens to his heart might well hear that heart-listening is not a reliable method of predicting the future accurately or forming true expectations, while the person who reasons inductively can never conclude on inductive grounds that induction is unreliable. This difference does not bear on the question we have asking, namely, whether inductive defences of the claim that induction is reliable are question-begging. The analogy shows what it is meant to show despite this difference. This difference between the cases might, though, prove important later, when we see whether it is possible to vindicate our preference for fixing our expectations inductively rather than by listening to our hearts.