



**OXFORD JOURNALS**  
OXFORD UNIVERSITY PRESS

---

Reliabilism, Induction and Scepticism

Author(s): David Papineau

Source: *The Philosophical Quarterly* (1950-), Jan., 1992, Vol. 42, No. 166 (Jan., 1992), pp. 1-20

Published by: Oxford University Press on behalf of the Scots Philosophical Association and the University of St. Andrews

Stable URL: <https://www.jstor.org/stable/2220445>

---

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



, and Oxford University Press are collaborating with JSTOR to digitize, preserve and extend access to *The Philosophical Quarterly* (1950-)

JSTOR

---

# *The Philosophical Quarterly*

---

## RELIABILISM, INDUCTION AND SCEPTICISM

BY DAVID PAPINEAU

This paper falls into three parts. In the first, I outline a defence of reliabilism that does not rest on intuitions about test cases. In the second, I address the sceptical problem of induction within this reliabilist framework. In the brief third section, I draw some general conclusions.

### I. RELIABILISM

#### *I.1. Knowledge and the project of enquiry*

Let me start with a question raised in Chapter 2 of Bernard Williams' *Descartes* (1978). Williams asks: why do human beings want knowledge? He takes it as given, as I shall throughout this paper, that humans want true beliefs. But, as we all know, a belief can be true and yet not be knowledge, as when it is a mere hunch or some other lucky chance. So the point of the question is: why do we want our beliefs to be *known*, in addition to being true?

Williams' answer goes as follows. Human beings are prone to false beliefs. So, if our desire for true beliefs is not to be idle, we will need to exercise ourselves to achieve it. It is no good, however, to start checking through all your beliefs with the intention of discarding the false ones. To have a belief is to take that belief to be true. So once you have formed your beliefs, internal inspection will not serve to distinguish the true from the false ones. In so far as you are prone to error, the damage will already have been done.

The only effective way for us to ensure that our beliefs are true is to block the error at source, by bringing it about that the processes by which we acquire beliefs in the first place are ones that generally yield

true beliefs. So Williams argues that the desire for true beliefs itself generates the desire that our beliefs should issue from processes that generally produce truths. And then, finding it independently plausible that beliefs produced by such processes should count as knowledge, Williams has an answer to his original question as to why we should want knowledge: our desire for knowledge derives from our desire to avoid error, in that attaining knowledge is the only effective *means* by which humans can avoid error.

I want to draw something more ambitious from this analysis. I think that, in addition to explaining why we should want knowledge, Williams' story also shows us what knowledge *is*. Williams takes it as given, from outside his analysis, that beliefs generated by truth-producing processes will count as knowledge. But I think that his story also explains *why* we have this concept of knowledge, why we pick out beliefs generated by a truth-producing process, *as knowledge*, as an especially good kind of belief. My idea here is that our concern to avoid error makes us especially interested in the state we need to get into as a *means* to avoiding error, and that this is why we call that state 'knowledge' – the state, to repeat, of having acquired a true belief from a process which generally produces true beliefs.<sup>1</sup>

### 1.2. *Certainty and reliability*

The above remarks prompt an immediate question: *how* truth-productive does a belief-forming process need to be in order to be an effective means of avoiding error, and therefore to qualify as a source of knowledge? In particular, it is enough that it merely be *reliable*, in the sense that it generally delivers true beliefs as a matter of contingent fact in this world? Or does it need in addition to yield *certainty*, in the sense that it should be *impossible* for a belief issuing from that process to be false?

Much traditional philosophical thinking assumes that knowledge requires certainty. But from the point of view of my remarks in the last section it is not clear why certainty should be necessary. Knowledge, I have suggested, is the state that we need to get into if we are to succeed in avoiding error. But we will succeed in this aim as long as we have

<sup>1</sup> Gettier cases call for the additional requirement that the truth of the belief not be an *accident* relative to the belief's deriving from the truth-producing process in question. This extra requirement is explained by my suggested hypothesis about knowledge, in that, if the truth of your belief is a Gettier-style accident, then the truth of your belief will not be a *result* of your being in the state someone concerned to avoid error wants to be in. See Williams (1978, pp. 43–4). Further analysis is needed, of course, to make the relevant notion of accident precise.

belief-forming processes which are reliable in this world. That such processes *would* lead us astray if things were different does not mean that they *will* lead us astray, as things are. This line of thought suggests that the traditional demand for certainty may be a mistake, perhaps fostered by an over-optimistic view of what human thought can achieve, but inessential to knowledge itself.

I shall return to the idea that the demand for certainty may be a mistake in section I.4, below. But first, in the rest of this section and the next, let me say a little more about the contrast I have drawn between reliability and certainty. Note that I have defined certainty objectively, rather than psychologically: the issue is whether it is in fact *impossible* for a given belief-forming process to produce a false belief, not whether the subject is *aware* of this, nor whether it yields some *feeling* of absolute security. It is this objective notion that matters to the arguments of this article. However, there are obvious links between it and subjective requirements on knowledge. For, as Descartes so forcefully argued, the only plausible source of certainty in the objective sense derives from various operations of the conscious mind, such as introspection and intuition. And so, if we can achieve knowledge with objective certainty, then we shall also, as it happens, have 'subjective warrants' available, in that we shall always be able to tell introspectively that our knowledge has come from these putatively infallible conscious sources.

From the reliabilist point of view, by contrast (henceforth I will use 'reliabilism' to mean the view that only reliability is required for knowledge), any subjective requirements on knowledge are gratuitous. For, in order for a belief-forming process to be reliable, there is no need for its reliability, or even its existence, to be available to consciousness. According to reliabilism, we will know, say, that there is a table in front of us, just in case the unconscious visual processes that give rise to such perceptual beliefs generally deliver true beliefs, whether or not we are aware of this. There is, therefore, no pressure, given reliabilism, to reconstruct such perceptual knowledge as first involving some infallible *introspection* of some sensory idea, and then some intuitively compelling *inference* from this idea to the presence of a table. If the demand for certainty in knowledge is unmotivated, then so too is this reconstruction of perceptual knowledge as involving infallible inferences from infallible introspections.

### I.3. *Knowledge and normativity*

It is sometimes felt that reliabilist epistemology changes the subject. As I have just pointed out, reliabilism implies that whether or not we know

will often hinge on matters, such as the reliability of some visual process, which lie quite outside our consciousness. But this seems to imply that we are at the mercy of nature, that we cannot *do* anything to affect whether or not we know. And this then makes reliabilist epistemology seem a quite different subject from the traditional version (henceforth I shall use 'traditional' to refer to views according to which knowledge requires certainty or subjective warrants<sup>2</sup>). For surely a central concern of traditional epistemology was the normative question of what we *should do* in order to ensure that our beliefs are knowledge.

However, this reaction to reliabilism involves a fallacy. It is true that traditional conceptions of knowledge offer advice on how to achieve knowledge: roughly, you should consciously monitor your thought processes, and avoid any which are not necessarily infallible. And it is true that reliabilism does not concur in this advice. But this is not because reliabilism has stopped offering advice on how to know, but simply because reliabilism offers *different* advice.

Where traditionalists advise aspirant knowers to monitor what goes on in their conscious minds, reliabilists will simply advise them to take whatever steps are needed to bring it about that their beliefs come from reliable processes. Such steps may well call for us to influence processes which lie outside consciousness, but that is no reason to conclude we cannot do this. After all, most of the things we influence lie outside consciousness, like our environments, our physical health, and so on. Similarly, there is no reason why we cannot influence non-conscious aspects of our belief-forming processes, by such means as rote learning, adjusting the working of instruments we rely on, and so on.<sup>3</sup> (Reliabilists will allow that conscious monitoring is *one* way to improve the reliability of our belief-forming processes. But it is not the only way.)

To guard against a possible misunderstanding, let me emphasize that I am *not* suggesting that it is a requirement on knowledge that knowers *must* take active steps to bring it about that their beliefs are knowledge. I know that I began this article by identifying the concept of knowledge as the state someone concerned to avoid error (a 'concerned enquirer' henceforth) wants to get into as a means of avoiding error. But it does not follow from this that the *only* way to be a

<sup>2</sup> What of the many philosophers, especially contemporary philosophers, who favour subjective warrants, but not certainty? I have two excuses for running the two requirements together. First, I can think of no good rationale, as opposed to intuitions, in favour of the demand for subjective warrants, except as a corollary of the desire for certainty. Second, most of my arguments will apply as much to the former demand taken separately as to the latter.

<sup>3</sup> For more on the extent to which we can alter our non-conscious belief-forming processes, see my *Reality and Representation* (1987, sections 7.4, 7.6).

knower is to take active steps to get into that state. For you may already be in the requisite state, not because you *did* anything to make your belief-forming processes reliable, but simply because they were reliable to start with. Such passive knowers will *already* be in the state concerned enquirers aim to get into, even though they do not themselves share the concern to get into that state.<sup>4</sup>

#### I.4. *Concepts, theories and intuitions*

I have just argued that reliabilism does not abandon the traditional normative issue of how best to acquire knowledge. In this section I want to consider a rather different argument for thinking that reliabilism changes the subject. This appeals not to considerations of normativity, but directly to intuitions about knowledge. Many philosophers take it to be *intuitively* obvious that subjective warrants are part of the concept of knowledge. And so they conclude that reliabilism, which dispenses with such requirements, must be wrong.

My response to this is that intuitions are not the only way to evaluate a theory of knowledge. There is a vast contemporary literature which aims to decide between reliabilism and traditional theories solely by appeal to intuitions about ingenious test cases (see Shope 1983). Unfortunately, however, these intuitions weigh on both sides, and the literature based on them is notoriously indecisive. By contrast, my approach in this article has not appealed to intuitions, but has tried to identify an underlying concept of knowledge, by locating the *role* it plays in our thinking, by trying to understand *why* knowledge is such a matter of concern to human beings.

My suggestion has been that knowledge is tied up with our desire to avoid false beliefs: it is the state a concerned enquirer needs to get into as a means of achieving this desire. Given this identification of the *concept* of knowledge, so to speak, we can then investigate more detailed conceptions, or *theories*, let us call them, of what that state is, more detailed theories of exactly what state an active enquirer needs to get into as a means to avoid error. The theory I am defending is that the

<sup>4</sup> Perhaps passive knowers should have at least this much in common with concerned enquirers: their belief-forming processes should not just *happen* to be reliable, but should be present *because* they are reliable. This suggestion will rule out purely fortuitous reliability, but will allow in reliability due to evolution, learning, and education, alongside the case where concerned enquirers adopt processes because they consciously recognize the reliability of those processes. This suggestion also indicates a possible answer to the delicate question of how to *individuate* belief-forming processes for purposes of assessing their reliability: namely, as a first approximation, we should individuate them by the same characteristics as are needed to explain their adoption. See Papineau (1987, pp. 136–8).

requisite state is acquiring a belief from a reliable process. The theory that certainty and hence subjective warrants are required is a *different* theory, but still a theory, in the terminology I am using, of the same *concept* of knowledge.

This is why I think that any intuitions which may favour the traditional theory over the reliabilist alternative are indecisive. If I have shown that the reliabilist theory is the *right* theory, in that acquiring a belief from a reliable process is indeed what is needed as the means to avoid error, then I have therewith shown that the traditional theory and the intuitions that support it are mistaken.

Perhaps there remains a gap here. Suppose it is granted that I have identified the underlying concept of knowledge successfully, and have shown that reliabilism is the theory of knowledge that best fits it. An obvious question which then arises is why anybody should have had contrary ideas about certainty and subjective warrants in the first place. For, as I have told the story, reliabilism, as opposed to the traditional theory, follows pretty quickly from the concept of knowledge, thus making it mysterious why anybody should ever have thought anything more was needed. We cannot simply rest with the suggestion, offered briefly in passing earlier, that traditional ideas about knowledge may have been fostered by excessive optimism about what can be achieved. For that we can get something does not explain why we should want it, if it is not already desirable.

I shall return to this issue at the end of the next section. But first let me comment briefly on the similarities between the strategy outlined in this section and that defended by Edward Craig in 'The Practical Explication of Knowledge' (1986). At a detailed level, Craig's views differ from mine: he offers a third-person account of the concept of knowledge, by contrast with my first-person account, arguing that knowledge is the state our *informants* need to be in for us to avoid error, not the state that *we* ourselves need to be in. But at the level of general strategy, Craig and I are in accord, in that he too seeks to offer an account of the *point* of the concept of knowledge, and to use this account to explain the nature of that concept, rather than simply trying to identify the concept from intuitions alone.

As to our differences, my objection to Craig's line would be that he is in effect focusing on the special case in which *we* succeed in avoiding error by acquiring beliefs from *informants* who succeed in avoiding error. I accept that this special case may well have been of primary significance in the historical development of the everyday concept of knowledge, in that worrying about your informants' reliability calls for rather less sophistication than worrying about your own reliability (see

Craig, 1986, p. 215). But, even so, Craig's third-person focus seems to me to have the disadvantage of cutting the link with the traditional normative issue of what *we* should do to avoid error. That is, even if it is unfaithful to the *history* of the concept of knowledge to view the desire for good informants as a special case of a general desire to have good belief-forming processes, I would argue that the more general perspective I have adopted nevertheless has the advantage of showing how the concept of knowledge relates to familiar *philosophical* worries about knowledge.

### I.5. *How reliable?*

*How much* reliability should a reliabilist require for knowledge? I shall consider two dimensions to this question. First, do we need 100 per cent reliability, or is some lesser per centage, such as 95 per cent, say, enough? It would be a mistake to think that the rejection of certainty has already decided this question in favour of something less than 100 per cent. For 'certainty', as I have been using it, implies that a belief-forming process *cannot* go wrong, will deliver 100 per cent true beliefs in *all* possible worlds. This is a much stronger requirement than 100 per cent reliability in this world. So we can reject certainty and still uphold a requirement of 100 per cent contingent reliability. On the other hand, even the latter seems a fairly strong requirement. So perhaps we should consider arguments in favour of some lesser per centage.

Second, over what range of possibilities is the per centage we settle on to be assessed? The argument so far has shown it would be a mistake to require this degree of reliability to hold up across *all* possible worlds. But, as we shall see, there may remain reasons for wanting it to hold up across at least *some* counterfactual situations.

Let me consider the per centage question first. Up to a point, it is possible to bypass this question. Suppose a given belief-forming process delivers beliefs which are true 95 per cent of the time. Then the appropriate output from that process would not be a *full* belief in the first place, but a 0.95 *degree* of belief. After all, if you believe it is going to rain tomorrow on the basis of a 95 per cent reliable method of forecasting, you would be ill-advised to bet a million pounds to a penny, or indeed to stake anything more than nineteen to win one, on its raining tomorrow. So if knowledge unqualifiedly requires *belief*, as I have implicitly been assuming throughout, then this in itself seems to call for belief-forming processes which deliver truths with 100 per cent reliability.

Still, perhaps it is rather quick to assume that knowledge requires strictly full belief. After all, in everyday discourse we certainly refer to beliefs of high, but less than strictly full, degree as 'beliefs' *simpliciter*, and to that extent we should expect the notion of knowledge also to encompass sufficiently well-founded beliefs of high, but not full, degree. But, having said this, there is then an obvious answer to the question of *how* well founded such a belief of high, but not full, degree needs to be to qualify as knowledge. For, under the present suggestion, everyday discourse has certain standards, perhaps varying from context to context, of how firmly a belief has to be held to qualify as a belief *simpliciter*. So why not simply incorporate those standards into our analysis of knowledge, and say that for a belief to be knowledge it should come from a process whose reliability is at least sufficient to warrant the degree of confidence required for the belief to qualify as a belief *simpliciter* in the first place, and not just as what even everyday discourse would consider as a partial belief?

It will be helpful for what follows to observe that, while it is certainly true that we often allow beliefs of less than strictly full degree to qualify as knowledge, there is also a practical sense in which it is always *better* to get a belief from 100 per cent reliable processes. The reason, to put it simply, is that such beliefs will then be true, and so decisions informed by them will succeed with probability 1, whereas, if those beliefs came from less than 100 per cent reliable processes, then the actions they informed would be less likely to succeed.

This is somewhat over-simple, it is true, in that even if your beliefs are proof against error in this way, they may still not be *informative* enough to tell you how to achieve some result; and, even if they are, you may fail to draw the inference correctly. But these two caveats would apply equally even if the same set of beliefs came from less than 100 per cent reliable processes, and so do not affect the point that it is always preferable, from the point of view of achieving your desires, to get a full belief from a 100 per cent reliable process rather than a less than strictly full belief from a less reliable process.

It should also be admitted that in many cases the extra *costs* of getting 100 per cent reliability will not be worth the extra probability of success, in which case we shall do better to settle for a partial belief. This is no doubt why everyday discourse does not make strictly full belief a precondition of knowledge in general. But this merely calls for a yet further qualification, and still does not affect the underlying point that, when costs are equal, full belief from 100 per cent reliable processes is always better.<sup>5</sup>

<sup>5</sup> I am here assuming beliefs about non-chance matters. Where genuine non-unitary

Let me now turn to the second question raised at the beginning of this section: what range of possibilities is relevant to the reliability of belief-forming processes? At first sight, it might seem to follow from my overall argument that reliability in the actual world is all that matters. After all, as I observed earlier, reliability in this world is all that we need in order to avoid error. However, there are good reasons why knowledge calls for more than merely this-worldly reliability. Let us return to the idea that knowledge is the state concerned enquirers need to get into in order to avoid error. It is true that concerned enquirers have no interest in reliability in non-actual worlds as such. Nevertheless, in acquiring processes which are reliable in this world, concerned enquirers will inevitably acquire processes which are reliable in a range of non-actual situations as well. The reason is that, if you are a concerned enquirer, you will not be able to anticipate the future in enough detail to be able to tell exactly when you are going to use any given belief-forming process, and so will not know exactly which truths that process needs to deliver in order to be reliable in the actual world. Instead, you will inevitably have only limited knowledge about the general nature of the world and your particular situation in it,<sup>6</sup> information which will narrow down the range of circumstances you may in future find yourself in, but certainly will not tell you exactly what they will be. So, in aiming for reliable belief-forming processes, you will inevitably be constrained to aim for belief-forming processes which will reliably deliver true beliefs across the entire *range* of possible circumstances that your current information leaves open for you. Since not all the possibilities in that range will become actual, you will inevitably be aiming to get into a state which would deliver true beliefs in various nearby possible worlds, as well as in the actual one.

So my overall approach to knowledge accommodates the requirement that knowledge should have a certain degree of counterfactual

chances are involved, the best degree of belief about any *outcome* will, of course, be different from 1. But even here it will always be better to get true beliefs of full degree about *chances*, belief-forming costs apart, for such beliefs will then ensure that in general your decisions *maximize* your objective chances of success.

<sup>6</sup> Some readers may be unhappy with these assumptions about the *knowledge* available to concerned enquirers. There are two possible worries here, one about a possible circularity on my part, the other about a possible circularity on the part of concerned enquirers. If you are worried about a circularity on *my* part, let me observe that my present concern is not to *define* knowledge, so much as to identify the *role* the concept plays in our thinking: I think it is helpful in *this* task to consider the predicament of a concerned enquirer who already has some knowledge; it would be a further task to try to specify, in *non-epistemological* terms, necessary and sufficient conditions for someone to be a knower. If, on the other hand, you are worried about some kind of circularity on the *concerned enquirer's* part, on the grounds, perhaps, that any seriously concerned enquirer ought to start by assuming nothing, then your worry will be addressed in section II below.

reliability. Note, however, that this is still a long way short of requiring reliability in *all* possible worlds, or even reliability in *all causally possible* worlds. For the information already possessed by concerned enquirers will still in general be enough to rule out the possibility of their being in most possible worlds (such as the world where you are manipulated by Descartes' evil demon) or even in most causally possible worlds (such as the world where you are a brain in a vat).<sup>7</sup>

I promised to return to the question of why certainty should be intuitively plausible as a requirement for knowledge. A serious answer to this question would include an historical dimension, examining the development of Western epistemological notions, with particular reference to the medieval distinction between demonstration and opinion, to the seventeenth- and eighteenth-century struggles to find a place for the newly emerging scientific knowledge within this distinction, and, perhaps most important of all, to the religious dimensions which so animated the participants in these debates. However, any such investigation is beyond the scope of this article. Instead, let me offer a possible conceptual diagnosis of the pull of certainty, not as a competitor to an historical account, but as a possible complement.

I have just argued that knowledge requires not only reliability in the actual world, but reliability across all worlds which are possible, relative to the information open to concerned enquirers. However, in discussing this issue of counterfactual reliability, I have so far implicitly been taking it for granted that our notional concerned enquirers are aiming for full beliefs from 100 per cent reliable processes. But, as we saw earlier, in many practical contexts it will often be more efficient to settle for less than full beliefs, delivered by belief-forming processes of appropriately high, but less than perfect, reliability. Now, an enquirer who was concerned to acquire such a less than full belief would be entitled to ignore, when assessing the reliability of the relevant belief-forming process, not only all worlds which are impossible relative to his or her current information, but also any worlds which fall below an

<sup>7</sup> Craig (1986, pp. 218–21) argues similarly from the limited informational situation of an enquirer to a counterfactual requirement on knowledge. But he is concerned with the kind of counterfactual reliability we want of our informants, given our limited information about the *particular* situation at hand, whereas I am interested in the counterfactual reliability we want of ourselves, given our limited *general* information about the situations we are going to be in.

The precise degree of counterfactual reliability required for knowledge is a complex issue, which I shall not pursue any further here, except to observe that this is an area where one good theoretical rationale seems to me to be worth a thousand delicate intuitions. Note that an appropriate counterfactual requirement ought to explain why the Gettier cases are not knowledge.

appropriate threshold of probability relative to that current information. For the fact that the process would be unreliable in such unlikely circumstances would not amount to a reason for the concerned believer not to attach a high, though less than full, degree of belief to its deliverances.

Now add in the consideration, elaborated earlier in this section, that, although it is often perfectly sensible to settle merely for a high degree of belief, it is always *better*, especially where it is important that your actions will not fail, to get full beliefs from 100 per cent reliable processes. The implication, then, is that your knowledge will get better, the more possible circumstances with any non-zero probability your belief-forming processes are reliable across. Now, I do not think that this line of thought provides a *good* rationale for thinking that the highest kind of knowledge demands reliability across *all* possible worlds, for, as I have said, most possible circumstances will be downright impossible relative to the information available to any concerned enquirer. But one can see how it would be easy to slide from the thought that you need ideally to guard against any possibilities that your information leaves with any non-zero probability, to the thought that you need ideally to guard against any possibilities whatsoever.

## II. THE PROBLEM OF INDUCTION

### II.1. *Problem? What problem?*

I turn now to the problem of induction. Let us suppose, for the sake of the argument, that the general form of induction is simple enumeration. (I do not really think this is a good model for inductive inferences, but it will help the philosophical exposition to assume so for a while.) So, for example, from the premise, that  $N$  ravens have been black so far, we conclude that *all* ravens are black. Schematically:

$$\begin{array}{l}
 (1) \quad Fa_1 \ \& \ Ga_1 \\
 \quad \quad \cdot \\
 \quad \quad \cdot \\
 \quad \quad \cdot \\
 \quad \quad Fa_N \ \& \ Ga_N \\
 \hline
 \quad \quad \text{All } Fs \text{ are } Gs
 \end{array}$$

The traditional complaint about this form of inference is that it is logically invalid. The conclusion does not follow logically from the premises. It is logically possible that the premises be true but the

conclusion be false. For reliabilists, however, this complaint has no force. Since a form of inference can well be contingently reliable without being logically guaranteed, reliabilists can simply respond to the traditional complaint by arguing that the illogicality of inductive inferences is no reason to deny that such inferences yield knowledge.

Perhaps it is worth pausing briefly to explain how the notion of reliability applies to inferences. Though I have not treated this explicitly so far, the appropriate notion is obviously *conditional* reliability: the conclusion should always be true in the actual world, *if* the premises are. (This will then ensure, for a reliabilist, that reliable inferences will transmit knowledge, that they will yield *known* conclusions when applied to *known* premises. For if the premises are known, in the sense that they are true and reliably arrived at, then any conclusion derived from a conditionally reliable inference will also be true and reliably arrived at, and so known.)

It is tempting to stop here, with the observation that the logical invalidity of induction does not mean its conclusions are not knowledge. However, I suspect that most readers will be unpersuaded by this quick way with inductive scepticism, even if they are persuaded by the general arguments for reliabilism. So in the rest of this section I shall consider whether there are any further reasons why a reliabilist should worry about induction. That is, from now on I shall take it as given that reliabilism is the right account of knowledge in general, and focus instead on the question of whether any sceptical doubts about induction can still arise *within* this assumption.

## II.2. *Is induction reliable?*

One possible worry about the simple reliabilist response to the problem of induction sketched in the last section is that it seems little different from the ‘analytic justification of induction’ proposed by Edwards (1949) and Strawson (1952, Ch. 9). Yet it is now widely agreed that inductive inferences cannot be shown to be legitimate simply by observing, as the ‘analytic justification’ does, that most people would characterize induction as a central case of ‘rational’ thinking. For such facts of common usage leave it open that there may be underlying requirements for a form of reasoning to be rational which are not in fact satisfied by induction, and that most people may therefore be in error in holding induction to be rational.

However, the reliabilist response to induction is quite distinct from the analytic justification. Reliabilists do not accept a form of reasoning as rational just because it is widely regarded as ‘rational’, but only in so

far as it satisfies the underlying requirement of reliably delivering truths.<sup>8</sup> In particular, reliabilists will deem induction to be rational, and its conclusions therefore knowledge, not because it is called 'rational', but because they believe that it is, in fact, a reliable method of getting new truths out of old ones.

However, this now points to an obvious problem. That induction reliably generates truths is itself a substantial contingent claim. Yet no support has so far been offered for this. We reliabilist friends of induction seem simply to be taking it for granted, in the context of our general reliabilism, that induction yields knowledge. But whether induction yields knowledge, given reliabilism, is supposed to be the point currently at issue.

Some reliabilists are inclined to respond, at this stage of the proceedings, that we do not need to *know that we know* in order to know.<sup>9</sup> I think this is the wrong move. It is perfectly true, of course, that ordinary non-philosophical knowers do not need to know that they know. But the present demand for a defence of the claim that induction yields knowledge is not being made of ordinary knowers who are *using* inductions, but rather of us philosophers who are talking *about* inductions, and in particular about the question of whether inductions yield knowledge. We reliabilist friends of induction are explicitly claiming that inductive inferences reliably yield truths. Given this, it is perfectly reasonable for someone to challenge us to provide support for this claim.

Of course, if we fail to meet this challenge, this will not necessarily show that induction does *not* yield knowledge. To lack any grounds for accepting the reliability of induction is not yet to have grounds for denying it. But such a stand-off would be a failure for us friends of induction, and a success for our sceptical challengers. The point at issue is whether induction yields knowledge, that is, given reliabilism, whether induction reliably generates truths. We friends of induction say yes; our sceptical challengers ask for support for this claim. If we cannot answer them, then they will have succeeded in showing we are not entitled to our stance.

<sup>8</sup> This claim perhaps deserves further discussion. Some reliabilists would hold that, while reliability suffices for *knowledge*, some kind of extra subjective warrant is needed for *rationality*. In my view, however, the arguments about knowledge go through for rationality. For a defence of the analogous point about 'justification', see Goldman (1979).

<sup>9</sup> See Van Cleve (1984, pp. 559, 562). Much of what follows in this section is influenced by Van Cleve's important article. However, in the passages referred to, Van Cleve seems to deny that reliabilists need to defend the reliability of induction, on the grounds that reliability is an 'external' requirement. This seems to me an unfortunate slip, given that evaluations of reliabilist defences of induction are highly sensitive to prior judgements on exactly what the reliabilist needs to do.

So to uphold induction as a source of knowledge we need to show that inductive inferences are reliable. But, now we are clear about this need, I do not think it is hard to satisfy. The obvious way to find out whether induction is reliable is to examine such evidence as bears on the matter. When people make inductions, do their conclusions turn out to be true? There are plenty of past examples of people making inductions. And when they have made inductions, their conclusions have indeed turned out true. So we have every reason to hold that, in general, inductive inferences yield truths. That is:

- (2) When person<sub>1</sub> induced, from  $N$  observations of  $A$  going with  $B$ , that 'All  $A$ s are  $B$ s', this conclusion<sub>1</sub> was true.

When person<sub>2</sub> induced, from  $N$  observations of  $C$  going with  $D$ , that 'All  $C$ s are  $D$ s', this conclusion<sub>2</sub> was true.

·  
·  
·

When person <sub>$N$</sub>  induced, from  $N$  observations of  $L$  going with  $M$ , that 'All  $L$ s are  $M$ s', this conclusion <sub>$N$</sub>  was true.

---

Whenever someone induces, his or her conclusion is true.

### II.3 *Varieties of circularity*

Let me first put to one side two obvious worries about the *premises* of this argument. First, are there not plenty of past examples of *unsuccessful* inductions with false conclusions, as well as successful ones with true conclusions? Second, how can we know that even the successful inductions *are* successful, given that observation of the past will only show, for example, that *As* have been *Bs so far*, not that *all As* are *Bs*? I shall deal with both these points in due course. But for the moment it will be helpful to ignore them, and attend instead to the move *from* the premises *to* the conclusion of (2). For it will not have escaped the notice of most readers that this is itself an inductive inference of just the kind whose reliability we are presently concerned to investigate.

However, is there anything wrong with this? It is not as if the discussion so far has identified some *flaw* in induction of a kind which would imply that it ought to be eliminated from the battery of procedures by which we normally arrive at our beliefs. In particular, we have agreed that the logical invalidity of inductive inferences in itself casts no discredit on induction. Given this, when a certain question of fact is raised – namely, are inductive inferences always

reliable? – what is more natural than to try to resolve this question by means of our normal procedures of investigation, which include, as it happens, our inductive procedures?

I know that to some philosophical sensibilities this will seem unduly complacent: surely we are not *entitled* to *any* methods of investigation until we have demonstrated their worth. But where are we supposed to start? We certainly need to begin with some methods of thought, lest we lapse into philosophical catatonia. Many philosophers, I realize, will want to follow Descartes and restrict our initial methods to introspection and intuition. But Descartes' rationale for this restriction was that it promised certainty, and we have already agreed that this is an unnecessary desideratum on our methods of thought. So why is it not as proper to proceed with our normal methods of thought, at least until we uncover some reason to distrust them? And this will include induction, since, to repeat, we have not as yet been given any reason to distrust induction.

Still, even if nothing has as yet been shown wrong with induction in general, it may well be felt that there is something wrong with the inductive argument (2) in particular. For is not (2) a *circular* argument, and therefore illegitimate? This objection needs to be treated carefully. It is true, as we shall shortly see, that circularity of a certain sort is present in (2). However, provided we keep firmly in mind the specific argumentative task to which (2) is directed, we shall also see that this circularity is not damning.

As a first step, we need to distinguish between '*premise circularity*' and '*rule circularity*'. An argument is premise-circular if its *conclusion is contained among its premises*; an argument is rule-circular if it reaches the conclusion that a certain rule of inference is reliable by *using* that self-same rule of inference (see Braithwaite 1953, pp. 276–7; Van Cleve 1984, p. 558). Clearly, premise circularity is a vice in an argument. The point of an argument is to take us from old beliefs, which we already accept as premises, to some new belief as a conclusion. But if the conclusion is already contained in the premises, then the argument will fail in this primary task. However, argument (2) is clearly not premise-circular. It is a genuinely expansive argument, whose conclusion, that all inductions yield true conclusions, manifestly outruns its premises, that *N* inductions so far have done so.<sup>10</sup>

<sup>10</sup> There is a problem about premise circularity. If 'containment' just means logical implication, then all deductive arguments will be premise-circular. Some philosophers, most notably, Descartes and Mill, take this to show that deduction is uninformative. The majority prefer to understand 'containment' more strictly. This debate is irrelevant to our current concerns, however, since both sides will agree that the non-deductive argument (2) is not premise-circular.

On the other hand, argument (2) *is* rule-circular. Even if the claim that induction is reliable does not appear among its premises, it does *use* an inductive inference to reach its conclusion that induction is reliable. I have a number of comments to make about the rule circularity of (2). But first let me make a wider comment, not about argument (2) in particular, but about rule circularity as such: namely, that it can scarcely be a general requirement, on all legitimate forms of inference, that it be possible to show that they are reliable in some non-rule-circular way. For this would disqualify even deduction as a legitimate form of inference. (While it is possible to demonstrate that deductive inferences are reliable – indeed, necessarily reliable – by means of the standard semantic soundness proofs, these demonstrations themselves unquestionably employ deduction.) So the fact that induction can only be shown reliable in a rule-circular way, as in (2), certainly does not in itself yield any immediate reason to distrust induction.

However, this is merely to repeat the point that we have as yet been given no good argument for *distrusting* induction. The issue we now need to address, however, is whether (2) takes us beyond this, and shows scepticism to be wrong, by giving us a good argument for *trusting* induction, despite the fact that (2) is admittedly rule-circular.<sup>11</sup> To see that (2) does succeed in this, let us recall the context of argument in which (2) was put forward. We agreed, on general reliabilist grounds, that induction does not need to be logically valid to yield knowledge, but will yield knowledge just in case it is reliable. However, the sceptic then pointed out, we cannot just take the belief that induction is reliable for granted. To which we responded that we are not taking this for granted, but have a good argument, based on empirical evidence, for the conclusion that induction is reliable, namely, argument (2). It seems to me that, in this specific context, the context of showing a sceptic who accepts reliabilism that we are not just helping ourselves to the belief that induction is reliable, (2) does just the job it is required to do.

Perhaps the best way of showing this is by detailing some of the tasks argument (2) *is not* intended to fulfil. For a start, we should recognize that argument (2) would be no good for *persuading* people who do not make inductions to start making them. While the *conclusion*, that inductions are reliable, would certainly be a good reason for such

<sup>11</sup> If it does, then this will answer the second of the questions raised about the premises of (2) at the beginning of this section, namely, the question as to how we know that the conclusions ('All *As* are *Bs*') of past inductions are true. The answer is, by induction. For, if the distinction between rule and premise circularity legitimates the inductive move from (2)'s premises to (2)'s conclusion, then it will also serve to legitimate antecedent inductive moves from instances of *As* being *Bs* to 'All *As* are *Bs*'. See Van Cleve (1984, pp. 560–1).

people to start inducing, if they accepted it, they clearly will not be persuaded to do so by (2), for the route from (2)'s premises to (2)'s conclusion requires just the kind of inductive inference that they eschew. In particular, then, (2) will be no good for persuading people who have already reflected on the reliability of induction, and have been persuaded, for whatever reasons, that they ought to stop performing inductions, that they ought to start again. However, in the present context of argument, this is no demerit in (2). Argument (2) is not addressed to people who avoid inductions. We may yet discover good reasons for avoiding inductions, and indeed in the next section I shall briefly examine some possible such reasons, but right now we are assuming that nothing has yet been shown wrong with induction, and are considering whether, given this, argument (2) can show us whether induction is reliable. So (2) should be thought of as addressed to people who have not yet been given any reason to distrust induction. And (2) ought surely to persuade them at least that induction is reliable.

I realise that many readers will feel that, if (2) is a good argument in defence of induction, then it ought to be capable of persuading *any* intelligent being, with whatever epistemological habits. But this is an extremely strong demand, and it is not at all clear why we should accept it. The only plausible rationale, once more, seems to stem from the assumption that knowledge requires certainty, together with the assumption that the only kinds of belief-forming processes which can plausibly deliver certainty are conscious operations whose logical infallibility is introspectively available. Together, these assumptions imply that any source of knowledge ought in principle to be recognizable as such by *any* conscious beings, in virtue of their introspective abilities, and thence that a good argument for the legitimacy of some source of knowledge ought to persuade any conscious beings, however wrong-headed their starting position. However, once we reject the assumption that knowledge requires certainty, then this line of reasoning falls away, and the strong demand that a good defence of induction ought to persuade any conscious being is left without any obvious means of support.

These last remarks bear on the question of 'counter-inductive' arguments for 'counter-induction'. It is often observed that inductive arguments for induction, like (2), have counter-inductive mirror images. Counter-inductivists, when they observe that a number of *As* are all *Bs*, conclude that the next *A* will *not* be a *B*. When it is pointed out that this is illogical, they can respond, 'So what? Illogical it may be, but this doesn't show that it is not in fact a good way of reliably reaching true conclusions.' And when we say, 'All right. But what basis do you

have for supposing that counter-induction *does* in fact deliver true conclusions?', they reply, 'Ah, so *that's* what you're worried about. Let us then look at the evidence that bears on the question. On a large number of occasions in the past people have counter-induced, and have been led to false conclusions. So we conclude – counter-inductively – that the next time we counter-induce we will get a true conclusion.'

There is room to dispute whether this is in fact a perfect mirror image of (2) (see Van Cleve, 1984, footnote 16). But let that pass. The more important point is that, even if counter-inductivists can mirror (2), this does nothing to discredit (2) itself. I have already conceded that (2) is not going to *persuade* people who do not make inductions to start making them. Counter-inductivism now simply gives us a further example of people who have abnormal inferential dispositions, and who will not therefore be persuaded by (2). Except that the parable of the counter-inductivists adds an extra twist, namely, that counter-inductivists *will* be persuaded, by their mirror of (2), to the conclusion that their abnormal counter-inductive dispositions are reliable. But all this leaves (2) untouched. We should not expect it to perform the impossible task of knocking imaginary non-inductivists out of their non-inductivism – its task is only to allow normal people, like ourselves, to resolve the issue of whether induction is reliable.

By this stage, some readers may be feeling that argument (2) does not do very much. Indeed, if it only works for people who already make inductions, is it really doing anything at all? My answer is that it is not *supposed* to do very much. Nearly all the serious work was finished before (2) came on the scene. Most importantly, the general arguments for reliabilism have already shown that the logical invalidity of induction is not a problem. Argument (2) is just supposed to show that, given that there is nothing problematic about induction, then there is no barrier to our concluding that it is reliable, and hence that it yields knowledge.

Trained philosophers naturally expect a 'justification of induction' to do something to rehabilitate induction, in response to an argument that there is something wrong with it. But (2) is not meant as a 'justification' in this sense. So we should not condemn it for its failure to be one.

### III. THE STRATEGY GENERALIZED

As a first step towards generalizing the anti-sceptical strategy of the last section, let us be more realistic about induction. I have already noted one way in which the above discussion of induction has involved an idealization, namely, in respect of the assumption that in our

experience all past enumerative inductions have been successful. This assumption is, of course, manifestly false. There are plenty of good examples of enumerative inductions leading to false conclusions, from Russell's chicken, who expected to be fed every day, to the Newtonian physicists who expected acceleration always to be inversely proportional to rest mass.

In any case, apart from such direct evidence, there is also a principled argument to show that simple enumerative induction cannot *possibly* be a reliable method of inference. I refer to Goodman's 'new problem of induction'. Goodman (1954) shows that there are far too many ways of classifying events, far too many *As* and *Bs*, for every instance of schema (1) to yield a true conclusion. Indeed, Goodman shows how to construct, for every instance of (1) that might yield a true conclusion, an infinity of other instances which will then yield false conclusions.

These are good arguments against enumerative induction. It is important to recognize, however, that they are quite independent of the *traditional* objection to induction. They do not just make the point that enumerative induction is logically invalid. On the contrary, they show that enumerative induction is not just invalid, but downright unreliable.

The moral, for us reliabilists who want to resist scepticism about induction, is that we had better not take our stand on simple enumerative induction as schematized in (1). At the very least, the form of induction we aim to defend will need some restriction on the *As* and *Bs* which are possible candidates for projection. Nor is this likely to be enough. It is implausible that all inductive unreliability can be laid at the door of non-projectable predicates. I suspect that, in order to identify a plausibly reliable form of induction, the conclusions of such inferences will need to be beliefs of less than full degree, and perhaps also restricted to claims that certain judgements are approximately, rather than precisely, true.

This is not the place to pursue such details. Instead, let me simply observe that the possibility of an alternative model of inductive inference opens the way to the anti-sceptical strategy outlined in section II once more. Imagine that we can show that our actual inductive practice is more sophisticated than simple enumerative induction, and that it cannot therefore be discredited as unreliable by either Goodman's new problem or past performance. And imagine, furthermore, that when we investigate its reliability, using existing methods of investigation, including those inductive methods, we find ourselves able to conclude that it is reliable. Then this defeats scepticism about

our inductive practice. As before, neither the fact that this practice may be invalid, nor the fact that its reliability might only be discoverable in a rule-circular way, will be a barrier to our concluding that it yields knowledge.

In this article I have focused on induction; but I think the same response to scepticism will work across the board. Take our belief-forming processes in general. If the only objection to them is that they are not certain, in the sense that it is *possible* that they should yield false beliefs, then this is no reason to believe that they are *not* reliable. And if, moreover, investigation shows that those methods *are* reliable,<sup>12</sup> then, any rule circularity notwithstanding, we shall be in a position to conclude that they yield knowledge.<sup>13</sup>

*King's College, London*

## REFERENCES

- Braithwaite, R. (1953) *Scientific Explanation* (Cambridge: Cambridge University Press).  
 Craig, E. (1986) 'The Practical Explication of Knowledge', *Proceedings of the Aristotelian Society* 87.  
 Edwards, P. (1949) 'Russell's Doubts about Induction', *Mind* 68.  
 Goldman, A. (1979) 'What is Justified Belief?', in G. Pappas (ed.), *Justification and Knowledge* (Dordrecht: Reidel).  
 Goodman, N. (1954) *Fact, Fiction and Forecast* (London: Athlone Press).  
 Papineau, D. (1987) *Reality and Representation* (Oxford: Basil Blackwell).  
 Shope, R. (1983) *The Analysis of Knowing: A Decade of Research* (Princeton: Princeton University Press).  
 Strawson, P. (1952) *Introduction to Logical Theory* (London: Methuen).  
 Van Cleve, J. (1984) 'Reliability, Justification, and the Problem of Induction', in French P., Uehling T. and Wettstein H. (eds), *Midwest Studies in Philosophy 10* (Minneapolis: University of Minnesota Press).  
 Williams, B. (1978) *Descartes* (London: Pelican).

<sup>12</sup> One of the aims of my *Reality and Representation* (1987) was to show how we might investigate all our standard methods of belief formation, as we might investigate any other natural phenomenon, and discover that those methods are by and large reliable.

<sup>13</sup> I would like to thank Barry Gower, Peter Lipton, Harvey Siegel, and the Editorial Chairman and referees of *The Philosophical Quarterly* for many helpful comments on an earlier draft of this article.