INDUCTIVE KNOWLEDGE

(For Routledge Companion to Epistemology)

Alexander Bird

1 Introduction

In this article I take a loose, functional approach to defining induction: Inductive forms of reasoning include those *prima facie* reasonable inference patterns that one finds in science and elsewhere that are not clearly deductive. Inductive inference is often taken to be reasoning from the observed to the unobserved. But that is incorrect, since the premises of inductive inferences may themselves be the results of prior inductions. A broader conception of inductive inference regards any ampliative inference as inductive, where an ampliative inference is one where the conclusion 'goes beyond' the premises. 'Goes beyond' may mean (i) 'not deducible from' or (ii) 'not entailed by'. Both of these are problematic. Regarding (i), some forms of reasoning might have a claim to be called 'inductive' because of their role in science, yet turn out to be deductive after all—for example eliminative induction (see below) or Aristotle's 'perfect induction' which is an inference to a generalization from knowledge of *every* one of its instances. Interpretation (ii) requires that the conclusions of scientific reasoning are always contingent propositions, since necessary propositions are entailed by any premises. But there are good reasons from metaphysics for thinking that many general propositions of scientific interest and known by inductive inference (e.g. "all water is H_2O") are necessarily true. Finally, both (i) and (ii) fail to take account of the fact that there are many ampliative forms of inference one would not want to call inductive, such as counter-induction (exemplified by the 'gambler's fallacy' that the longer a roulette wheel has come up red the more likely it is to come up black on the next roll). Brian Skyrms (1999) provides a useful survey of the issues involved in defining what is meant by 'inductive argument'.

Inductive knowledge will be the outcome of a successful inductive inference. But much discussion of induction concerns the theory of confirmation, which seeks to answer the question, "when and to what degree does evidence support an hypothesis?" Usually, this is understood in an *incremental* sense and in a way that relates to the rational credibility of a hypothesis: "when and by how much does *e* add to the credibility of *h*?", although 'confirms' is sometimes used in an absolute sense to indicate total support that exceeds some suitably high threshold. Important but largely unanswered questions relate these topics, for example "does inductive inference correspond to the case of absolute confirmation for some suitable threshold?" I shall discuss inference and confirmation together, though it should be noted that some approaches eschew inference altogether. For example, the Bayesian takes scientific reasoning to be a matter of adjusting credences in propositions in the light of evidence, and says nothing about unqualified belief in a proposition. However, if we are interested in inductive *knowledge* then we must consider inference, since only then do we have a detached proposition that is the possible content of a mental state of knowing.

2 Enumerative induction

The form of inference discussed above and sometimes called simply 'induction' is a matter of inferring from a sample to the whole of a population. In the paradigm of enumerative induction one argues, as in the examples concerning planetary orbits, as follows:

(E1) all Fs in $P' \subset P$ are Gs
therefore
all Fs in P are Gs.

Where (E1) articulates a rule of inference, the corresponding notion of confirmation is *Nicod's criterion*: an F in P that is also a G confirms the generalization that all Fs in P are Gs. My characterization of enumerative induction, as an inference from a sample to all of a population, is more general than (E1) in order that it should encompass also:

(E2) the proportion of Fs in $P' \subset P$ that are Gs is p
therefore
the proportion of Fs in P that are Gs is p.

(E1) is a special case of (E2) where p=1, but (E2) cannot be considered as a matter of generalizing facts about individual members of P′. Rather (E2) concerns population level facts.

(E1) is most familiar in the form (E1′) 'all known Fs are Gs *therefore* all Fs are Gs'. Another popular way of expressing enumerative induction is (E1′′) 'all known Fs are Gs *therefore* the next F to be examined will be G', which we can gain from (E1) by putting P′ = the known Fs and P = the known Fs plus the next F to be examined. A generalization of (E1′′) yields a familiar version of (E2), Reichenbach's 'straight rule' of induction:

(E3) of n known Fs, m are G
therefore
the probability that some unknown F (e.g. the next F to be examined) is G is m/n

Just as (E2) is a form of enumerative induction, although not a generalization of facts about individuals, we should consider as inductive in this sense various other statistical inferences such as:

(E4) the mean value of parameter L in $P' \subset P$ is μ
therefore
the mean value of parameter L in P is μ.

In general the various techniques of classical statistical inference should be seen as refined instances of enumerative induction; in particular classical statistical inference seeks to correct the defects of (E1), *viz.* that it (a) does not tell us how large P′ needs to be before we can make the inference concerning all of P, and (b) it does not tell us how confident we should be in the conclusion when we draw it.

The principal weakness of the various forms of enumerative induction is that their scope is severely limited. Note that the conclusions in (E1)–(E4) mention only properties and parameters already mentioned in their premises. Therefore enumerative induction is unable to yield knowledge concerning entities of a kind of which we do not already know. Yet this is a crucial part of science—witness our beliefs in subatomic entities and their properties, or in plate tectonics, or stellar evolution, and so forth.

As Reichenbach's straight rule exemplifies, it is natural to seek to relate confirmation to probability. This thought lies behind Carnap's inductive logic. According to Carnap, inductive logic should be seen as a generalization of deductive logic where the conclusions are drawn only with a certain degree of probability. The degree to which evidence provides absolute confirmation for a hypothesis is the same as the probability of the hypothesis given that evidence: $C(h, e) = P(h|e)$, where the conception of probability being used is a logical one. The latter operates as follows. Consider a language with predicates denoting properties and names denoting objects. One can construct a complete 'state description', a maximal description of a way things can be, by saying of each object whether each predicate or its negation holds of it. The simplest approach to logical probability would ascribe to each state description the same probability. The conditional probability $P(h|e)$ is now fixed and so therefore is our confirmation relation. The drawback of this approach is that it does not allow any room for inductive learning. One might have thought that repeated observations of Fs that are Gs (without non-G Fs) would raise the probability that the next F to be examined will be G. But this simple approach to inductive logic does not yield that outcome. Carnap's important move is to concentrate not on state descriptions but on *structure* descriptions. '$F\mathbf{a}\wedge\neg G\mathbf{a}\wedge\neg F\mathbf{b}\wedge G\mathbf{b}$' is a different state description from '$\neg F\mathbf{a}\wedge G\mathbf{a}\wedge F\mathbf{b}\wedge\neg G\mathbf{b}$'. But they both have the same structure: 'one thing is F but not G and the other thing is G but not F'. Carnap now distributes probabilities equally among structure descriptions, rather than across state descriptions; then the probability assigned to a structure description is divided equally among the state descriptions with that structure. This distribution of probabilities, m^* yields a confirmation relation c^* that *does* allow for learning from evidence. A central problem for such an approach is to articulate why the distribution m^* is *a priori* more suitable than some other distribution (e.g. the simple distribution, $m^\dagger$, that gives each state description the same probability). If m^* isn't the *a priori* right distribution, then in what sense is this account or probability (and confirmation) *logical*?

3 Hypothetico-deductivism

Note that in (E1) the conclusion of the inductive inference entails the premise. In (E2) and (E3) the conclusions make the premises likely (without entailing them). According to (E1)–(E3), inductive support occurs where the inductive conclusion deductively entails the evidence or makes it likely. Hypothetico-deductivism takes this as the central idea in confirmation. Thus:

(HD) *e* confirms *h* iff *h* entails *e*.

Here h may be considered to be, for example, the combination of a theory and a set of auxiliary hypotheses. Also plausible, but less frequently discussed is the more general:

(HD′) e confirms h iff h entails that e is likely.

The advantage of hypothetico-deductivism over enumerative induction is that its scope is much wider. Enumerative induction (E1) can be seen as a special case of (HD) for the case where h is a generalization of e; Nicod's criterion is a consequence of (HD). But hypotheses can entail evidence without being generalizations of the evidence. Thus Fresnel's wave theory of light entails that there is a bright spot at the centre of the shadow cast by a disk, and is confirmed by the observation of such a spot, even though the wave theory concerns unobservable features of the world and so is no kind of generalization of the observational evidence in question.

This advantage turns out, however, to be a disadvantage, when we see that the deductive relationship allows for too liberal an account of confirmation. The most famous of such problems is the Ravens Paradox. The hypothesis that all ravens are black, combined with the auxiliary hypothesis that object x is not black, entails that x is not a raven. So observing x, a white shoe, provides confirmation for the hypothesis that all ravens are black. Nicod's criterion delivers this conclusion since the white shoe confirms the hypothesis that all non-black items are non-ravens, which is logically equivalent to the hypothesis under test. Many find the conclusion absurd, though others accept it, merely regarding the support as very weak.

There is another problem. Hempel's *special consequence* principle tells us that when e confirms h, then e confirms any consequence of h. This seems reasonable: recall (E1′′) 'all known Fs are Gs *therefore* the next F to be examined will be G', which I claimed to be a special case of (E1), for the population consisting of the known Fs plus the next F to be examined. But strictly (E1)′′ follows only if we take the confirmation of the hypothesis that the known Fs plus the next F are all G to entail the confirmation of the proposition deduced from it, that the next F is G. Now assume for the following that h entails e. Therefore $h \wedge p$ also entails e, for any p. According to (HD) not only is h confirmed by e, but also h', where h' is $h \wedge p$. The special consequence condition tells us that since e confirms h' , e confirms any consequence of h', and so, in particular e confirms p. But p was an arbitrary proposition. So any proposition can confirm any other (putting $e=h$ makes that especially clear). It seems obvious that the special consequence principle is at fault. It can only be unrestrictedly true in the case of *absolute* confirmation. However, some restricted version of the special consequence principle would appear to be true for incremental confirmation, if there is ever to be any nontrivial ampliative confirmation—as, for example, in (E1′′). One would expect one's theory of confirmation to provide an answer: when evidence confirms an hypothesis, which logical parts of the hypothesis get confirmed and which do not? One response would be think that while the conjunction $h \wedge p$ may get confirmed as a whole, this is due just to the confirmation of h whereas p itself gets no confirmation, and this is because p plays no role in deducing e. But now consider $p \wedge (p \rightarrow h)$. The proposition e cannot be deduced from the second conjunct alone, so now p does play a deductively essential role—even though $p \wedge (p \rightarrow h)$ is logically equivalent to $h \wedge p$. As a theory of confirmation (HD) is thus incomplete; moreover, it does not point towards any obvious satisfactory supplementation.

4 Grue—the New Riddle of Induction

A feature of (E1) and (HD) is that they suggest that inductive confirmation is a *formal* relation. In deductive logic, the relevant relations, such as deducibility hold in virtue of the syntactic form of the relevant propositions. According to hypothetico-deductivism, the same is true of confirmation, because confirmation is held to be the converse of deduction. Likewise enumerative induction holds inductive inference to be licensed on grounds of the formal relation of conclusion to evidence, the latter being just a generalization of the former and generalization being a formal operation.

A major challenge to both these accounts of induction comes from Goodman's (1954) 'New Riddle of Induction', which shows that confirmation cannot be a formal relation. Define the predicate 'grue$_t$' as holding of *x* precisely when *x* is green and is first observed before time *t* or *x* is blue and is not first observed before time *t*. Consider the hypothesis 'all emeralds are green'. Given the auxiliary proposition that emerald *a* is observed before time *t*, then we may deduce that *a* is green. According to (HD) therefore, the observation of a green emerald before *t* provides confirmation of the hypothesis that all emeralds are grue. But the latter hypothesis has as a consequence the proposition that emeralds first observed after *t* will be blue. And this surely is *disconfirmed* by the observation of a green emerald before *t*.

This latter result may be seen as a particular case of the claim made in the preceding section that hypothetico-deductivism makes confirmation of hypotheses too easy. The same goes, however, for the more restrictive enumerative model of induction. Consider (E1) where 'F'='emerald', 'G'='grue thing', P′ is the set of emeralds observed to date (which is before *t)*, and P is the set of all emeralds (or is some subset thereof that includes emeralds first observed after *t* or never observed at all). According to (E1) we are entitled to infer that emeralds first observed after *t*, or never observed, are grue, and hence, by the definition of 'grue', are blue.

Neither of these results is acceptable, and so we should conclude that inductive confirmation is not a formal relation. Consequently we should not think that inductive relations are like deductive relations, which can be formalized in an *a priori* logic. The same applies to Carnap's inductive logic since the probability distribution *m** is relative to a language, with the consequence that that confirmation is also relative to a language. This means that his inductive logic fails to be logical in a key sense. A satisfactory deductive system should have the feature that if two sentences are deductively related in one language, their translations into another language should also be deductively related in the same way. But that does not hold for Carnapian confirmation.

5 Abduction and inference to the only explanation

The hypothetico-deductive model of confirmation is closely related to the deductive-nomological model of explanation, according to which:

(DN) Laws L and conditions C together explain fact *f* iff L&C deductively entail *f*.

Consider (HD). If our hypothesis is that laws L hold and conditions C obtain (i.e. *h*=L&C), then *e* confirms this hypothesis precisely when, according to (DN), L&C would explain *e*. Just as the hypothetico-deductive model of confirmation suffers from problems, so does the deductive-nomological model of explanation. In particular, the relation of explanation is not a formal one, *contra* (DN). From the law that anyone who ingests a pound of arsenic will die within 24hrs, and the fact that Jones has ingested a pound of arsenic, we may deduce the fact that he died within 24hrs. But in fact this does not explain his death because he was killed by being hit by a bus before the arsenic could kill him (Achinstein 1983). So background knowledge is required in addition to a deduction in order to work out whether we have an explanation. (Achinstein's case may also be regarded as a counter-instance to (HD). Let us say that we are investigating the hypothesis that a pound of arsenic will kill within 24hrs. The discovery that Jones dies within 24hrs of eating a pound of arsenic confirms that hypothesis as (HD) says it should. However, when we learn that Jones was in fact killed by a bus, the fact of his death now lends our hypothesis no support at all. Nonetheless, the death is still deducible from the hypothesis plus the fact of his eating the arsenic. One way to understand this is to consider that the deducibility relation is *monotonic*—if it holds at all, no further addition of information will prevent it from holding. Thus the confirmation relation should also be monotonic. But it is not, as this case shows.)

Even if both (HD) and (DN) are mistaken, what may be correct is the relation between them:

(A) *e* confirms *h* iff *h* would, if true, explain *e*.

The left to right implication in (A) will not do, however. Incremental confirmation is typically (if not always) transmitted by deduction, whereas explanatory power is not. For example, rising sea levels may be explained by, and confirm the hypothesis of, global warming. From the latter we may be able to deduce that there will be increased droughts, in which case the observation of rising sea levels may confirm that hypothesis that there will be increased droughts. But those future droughts do not explain current increases in sea levels.

The right to left implication in (A) remains plausible, nonetheless. Abductivist (or explanationist) conceptions of confirmation take explanation to be central to inductive confirmation. Note that abductivism does *not* say that *e* confirms *h* if *h* in fact explains *e*, since to know that latter would require already knowing that *h* is in fact true. Rather, abductivism says that the fact that *h* is a *potential* explanation of *e* provides confirmation to *h*.

Abductivism is able to encompass enumerative induction if one thinks that in the cases where enumerative induction lends confirmation, that is because some relevant fact provides explanatory power. For example, observations of planetary motions confirm Newton's law of gravitation via enumerative induction because the law explains its instances (Armstrong 1983; Foster 1983). In other cases, a common cause explains the known correlation and confirms its extrapolation; we can extrapolate the correlation between a high barometer reading and bad weather, because high pressure is a cause of both. The hypothesis that all emeralds are grue is not confirmed by the observation of green emeralds since background knowledge tells us there is no putative law or causal connection covering all the instances.

Abductivism, the claim that being a potential explanation of some evidence confirms a hypothesis, like (HD), fails to tell us *how much* this explanatory relation supports the hypothesis and likewise when we may infer that the hypothesis is true. Inference to the best explanation (IBE) aims to provide a more detailed account of abductivism. IBE employs the intuitive idea that some hypotheses are better potential explanations of the evidence than others, and regards better explanations as more likely to be true. To be more precise, IBE holds that under certain conditions, it is reasonable to infer that the best of a set of competing explanations is the actual explanation and hence is true. The conditions are: (i) that the best explanation is clearly better than its next best competitor, and (ii) that the best explanation is good enough (it meets some threshold of goodness—if it does not, then we may suspect that the problem of *underconsideration* applies). A key issue for any thorough account of IBE is to explain what explanatory goodness is.

Lipton characterizes IBE as a two stage process. In the first stage the imaginative capacity of the scientist generates a set of possible explanations of a phenomenon. In the second stage the generated hypotheses are ranked according to their explanatory goodness, and the top ranked is selected.

Three problems face IBE (Lipton 2004: 142–51). The first stage encounters the problem of *underconsideration*. The ranking by goodness at stage two cannot be any guide to truth if the actual (true) explanation is not among those considered by the imaginative power of the scientist at stage one. Even if the actual explanation is among those considered, stage two raises the two remaining problems, which Lipton names *Hungerford's objection* and *Voltaire's objection*. The former notes that beauty is in the eye of the beholder; that is, explanatory goodness is too subjective a quality to be correlated with objective truth. Even if goodness is objective, Voltaire's problem asks why it should be correlated with the truth. Presumably there are possible worlds where explanations that we would judge to be poor explanations are in fact very often true. IBE therefore assumes that ours is the best possible world, explanation-wise. Why should we think that this assumption is correct? Lipton notes that if the second stage ranking is accurate that shows that underconsideration cannot in fact be a problem, since ranking is a theory-laden process, and the reliability of the ranking implies the truth of the relevant background theories. In particular, background theories also play a role in setting our standards of loveliness. And so, successful inferential practices will be virtuously reinforcing. (In response to Hungerford's objection, Lipton says that while loveliness might be audience-relative, so also is inference.)

6 Bayesian confirmation

Bayesian epistemology avoids many of the problems facing other accounts of confirmation, including Hume's problem (see below). It does this by focusing, in its standard subjective form, on rationality of incremental changes to credences. Bayes's theorem:

$$\text{(B)} \qquad P(h \mid e) = \frac{P(e \mid h)P(h)}{P(e)}$$

is derivable from the standard axioms of probability and so is *a priori*. Subjective Bayesianism tells us that the probabilities in question are subjective degrees of belief (credences) and that if one receives evidence e then one's credence in h should now be made equal to $P(h|e)$ as given by (B), known as Bayesian conditionalization. One's old credence in h is multiplied by $P(e|h)/P(e)$ (where $P(e)$ is one's credence in the evidence, and $P(e|h)$ is once credence in the evidence, given the hypothesis).

An interesting question concerns the relationship between Bayesianism and other *prima facie* conceptions of inductive confirmation. Consider, for example, the case where the evidence is deducible from the hypothesis, as in the hypothetico-deductive model of confirmation. Then the old credence will be multiplied by $1/P(e)$, which so long as the evidence proposition is previously unknown, will be greater than one. In this way Bayesianism can encompass (HD) and does better both by giving an account of (HD′) and by giving a quantitative measure of incremental confirmation in these cases (the more unexpected the evidence, the better evidence it is). Bas van Fraassen argues that Bayesian confirmation and IBE are in conflict, and that because it can be shown to be irrational to conditionalize in a way that diverges from the Bayesian prescription (thanks to so-called Dutch book arguments), it follows that IBE is in error. Lipton, on the other hand, regards Bayesianism and IBE as compatible. He regards the explanationst considerations that are brought to bear by IBE as heuristics that guide our estimation of $P(e|h)$.

Since Bayesianism deals in the rationality of incremental changes in subjective probabilities, it makes no claim to say what probabilities one's beliefs should have. It tells one only what one's new probabilities should be, once the evidence has been received, given one's old probabilities. Except insofar as those old probabilities were themselves based on evidence subject to Bayesian conditionalization, those old are rationally unconstrained. Two people who have the same evidence will find that conditionalization brings their credences closer together. But if they start off with sufficiently divergent distributions of probabilities they will end up with credences that remain very far apart. Since Bayesianism fails to tell one what one should believe nor even how much one should believe certain hypotheses, it is unsuited to giving us an account of inference and so of inductive knowledge. Thus it evades Hume's problem principally by limiting its ambitions. For these reasons, it cannot be that IBE functions *merely* to estimate $P(e|h)$. IBE tells us about inference—what one may believe in the light of the evidence. (For a detailed exposition of Bayesian epistemology, see the chapter of that name in this volume.)

7 Hume's problem and the reliabilist response

The best-known and fundamental problem with any account of the capacity of inductive reasoning to yield knowledge is Hume's problem. What Hume's actual intentions were with this problem are a matter of debate (see "Hume" in this volume), but its basic structure is clear. When using some form of reasoning, it seems appropriate to ask whether use of that form of reasoning is justified. Prospects for an *a priori* justification are poor for those forms of inductive reasoning, the majority, that are ampliative. The very obvious efficacy of inductive reasoning, in science and everyday thought, would seem to be an obvious source of justification. However, the inference from the fact that induction has worked successfully for us on many occasions to its general reliability (we can expect induction to be reliable in current and future applications) is itself an inductive inference, an instance of (E1). And so this attempt to justify our inductive practices itself employs an inductive form of inference. As such this justification is circular and so may be thought to fail. Inductive reasoning thus seems to be without justification, and so the products of inductive reasoning cannot be knowledge.

One response to Hume's problem is to regard it as decisive. In which case, if one regards science as rational, one must propose a non-inductive basis for scientific reasoning. Sir Karl Popper's falsificationism attempts to do precisely that. Popper (1959) advocated hypothetico-deductivism, but regarded only one special case as admissible, that of absolute disconfirmation: the case where from a hypothesis *h* one deduces consequence *c*; one observes that *c* is false; hence one infers that *h* is false. Falsificationism is a highly sceptical view: although general hypotheses can be known to be false, they cannot be never known to be true.

Alternatively, in order to avoid scepticism about science, one may seek a diagnosis of Hume's problem that allows us to reject its conclusion. Hume's problem assumes that for S's use of a method M to produce knowledge, it must be the case that M is justified in a manner that is accessible to S. This requirement is a manifestation of epistemological internalism. The *internalist* rejects while the *externalist* accepts the idea that the ability of a process to justify beliefs may depend on some feature of the world that is not accessible to the user of that process (see "Internalism/externalism" in this volume). Goodman's new riddle raises similar problems. That showed that two inductive arguments may have the same syntactic form but differ in their confirmatory power. So confirmation depends on semantic features, such as whether the predicates in question denote natural properties and kinds or not. Those semantic facts are external to the structure of an inductive argument and the truth of its premises. They may be known to the subject, but only as part of background knowledge acquired by science, i.e. by some prior piece of inductive reasoning. Such knowledge will not be *a priori*.

Reliabilism is one, natural implementation of externalism that offers an explanation of the possibility of inductive knowledge: for a belief to count as knowledge it must be produced by a reliable process or method, one which produced true beliefs in an appropriate range of counterfactual circumstances (Armstrong 1973, Nozick 1981; see "Reliabilism" in this volume). If the world is in fact non-accidentally regular (e.g. law-governed) in certain respects, then enumerative induction, or certain classes of enumerative inductive inferences at least, may be reliable, and so knowledge generating. Since we are concerned with inductive *inference*, reliability will be a matter of the inferential rule (or pattern) producing true beliefs when given true beliefs as premises. A related reliabilist account may determine when inductively formed beliefs are *justified* (indeed reliabilism is typically seen more as an account of justification).

The reliabilist holds that in order for an inductively inferred belief to be knowledge (or to be justified) it must meet two kinds of condition (i) an *evidential* condition, that the premises of the inference are suitable, and (ii) a *reliability* condition, that the inference rule employed is a reliable one. As an externalist account, this does not in general require that the subject be aware that the reliability condition holds. Reliabilists and externalists more generally hold that for an inductive belief to be justified, it *can* suffice that the reliability condition in fact holds (in addition to the evidential condition). Thus Mellor (1991) argues that inductive habits can yield warranted beliefs thanks to natural, contingent, regularities, independently of one's knowing that one has such a warrant. However, one might reasonably ask whether, as a matter of fact, S can have, in addition to her warranted (justified) belief in *h*, a justified belief in the reliability of R, the rule that led to her belief in *h*. We do think we have this sort of knowledge and justification. When S's experience in baking leads her to believe a certain technique will lead to a firm crust we believe that her belief is justified. But when S* uses tea-leaves to predict the future, we believe that S*'s resulting beliefs will not be justified. Can our beliefs about the methods of S and S* be justified? And can S herself have a justified belief that her inductively inferred belief (in the future outcome of her baking) is correct?

Let us say that S's rule of enumerative induction, R, is indeed reliable, so that S's belief in *h* is justified. We may be able to show that R is reliable, for example by finding that R delivers true beliefs (e.g. predictions about the future that are subsequently verified) in a large number and variety of cases, while delivering no false beliefs (or only a few). If our rule of reasoning in this case is reliable, then by the reliabilist view of justification, our belief that R is reliable is itself justified. And S herself can engage in this reasoning also and acquire the same justified belief. Now let us consider the rule of reasoning, R*, just used to establish the reliability of R. It is clearly a form of enumerative induction. If R is itself a general rule of enumerative induction, then R* may be identical to R. Thus we will be using an inductive rule to establish the reliability (and so justify our belief in the reliability) of that very same rule. This would appear to be the very circularity that Hume warns us vitiates any attempt to justify induction.

The standard reliabilist reply draws upon the distinction between *premise-circularity* and *rule-circularity* (Braithwaite 1953: 25592; see also van Cleve 1984). The former is the kind of circularity in reasoning that is vicious and with which we are familiar—seeking to establish the truth of proposition Q using an argument among whose premises is the proposition Q itself. Rule circularity, however, is something different, and arises when one employs an argument to establish some proposition concerning a rule R, e.g. that it is reliable, and that argument-form is an instance of that same rule R. The key difference is that in rule-circularity there is no premise asserting the reliability of R, and so the conclusion is *not* among the premises. The Humean critic might object that by employing the rule R we are implicitly assuming the reliability of R, and so after all we are assuming what we set out to prove. The reliabilist (and, more generally, externalist) response is that a subject need not have any tacit belief or assumption concerning the reliability of the rules she in fact uses. Normally, it will suffice for knowledge or justification that the rule is in fact reliable. And so while premise-circularity does undermine the epistemic value of the conclusions, rule-circularity does not. Reliabilism itself faces many questions (see the chapter on Reliabilism for more details), for example: What degree of reliability is required for justification, and does this differ from the degree of reliability required for knowledge? How do we delineate the rule or belief-forming method that a subject is using? (One and the same particular inference may be considered as an instance of many different inference rules, some of which may be reliable while others are unreliable. Which is *the* rule employed?) Thinking about inductive inferences, is there one general rule 'enumerative induction' that has the form of (e.g.) (E3), the straight rule of induction? Or are there many rules, methods, and habits that are inductive in character but which are specific to certain subject matters or circumstances? (If so, there may not even be rule-circularity in the reasoning process described above. Note also that the details of the reliability condition required for inductive knowledge will be specific and differ from one inductive practice to another.) These questions are linked, because one might wonder whether one very general rule of induction is likely to be reliable enough to give us inductive knowledge—after all we frequently do get false beliefs from induction; a more specific rule can have better prospects of high reliability.

Hume's problem affects not only enumerative induction but other kinds of ampliative inductive inference such as IBE. The response to the problem of underconsideration and Voltaire's objection in effect depends upon its in fact being the case that our capacity to generate hypotheses does tend to succeed in including the actual explanation and its also being the case that our standards of loveliness match the way the world is. If asked to justify those claims, one might suggest that the best explanation of our success in employing IBE is that our hypothesis-generating and loveliness-judging capacities are effective in these respects. But this explanationist argument is just to use IBE in the justification of IBE, exhibiting the very circularity that Hume identified. A reliabilist answer may be given again here. So long as these capacities are indeed effective, IBE will be a reliable way of generating first-order beliefs, which will thus be justified. Furthermore, the second-order argument that IBE is indeed reliable can, for the same reasons, give a justified belief in IBE's reliability. The circularity identified is again only rule-circularity, not premise circularity.

8 Eliminative induction

One might worry, nonetheless, that IBE is insufficiently reliable to generate inductive *knowledge*. Are our judgments of loveliness so good that the hypotheses judged loveliness are always true? As Laudan (in effect) notes, many favoured, lovely hypotheses have been found to be false. Note that the ranking process implicitly assumes that the lower ranked hypotheses are still all consistent with the evidence. Perhaps IBE can only justify an ordering of our credences, not an inference that the best is in fact true.

A natural extension of this thought would suggest that we can only acquire knowledge of the truth of a hypothesis when inferred from evidence, if that evidence serves to *eliminate* competing hypotheses. *Eliminative induction* goes back at least to Francis Bacon but has been supported by a number of contemporary philosophers (e.g. Earman 1992; Papineau 1993; Kitcher 1993; Bird

2005). Considered as a deductive inference:

(L1) one of hypotheses $h_1, \ldots, h_i$ is true; hypotheses $h_1, \ldots, h_{i-1}$ are false; *therefore* hypothesis h_i is true

The capacity of eliminative induction to deliver knowledge of its conclusion depends on our ability to know that the first premise is true, i.e. to know of some suitably limited range of hypotheses, that the true hypothesis is among them. 'Suitably limited' here means sufficiently limited that it is possible to know the second premise; that is, sufficiently limited for us to falsify all but one of the hypotheses referred to in the first premise. Another approach would be to cast eliminative induction as non-deductive:

(L2) hypotheses $h_1, \ldots, h_{i-1}$ are false; *therefore* hypothesis h_i is true.

In this case, the inference is reliable when the subject has a reliable disposition to infer the truth of *hi* from the falsity of $h_1, \ldots, h_{i-1}$. Here one appeals to reliabilism again: to generate knowledge it may suffice that this inferential disposition is in fact reliable, whether or not the subject knows this. And in this case the disposition may be sufficiently specific that it is indeed sufficiently reliable. Alternatively, one might try to argue that the premises of (L1) can be known in certain cases. For example when the hypotheses are explanatory, it may be possible to know that the possible explanations of some phenomenon are limited to a constrained set of hypotheses:

(L3) only hypotheses $h_1, \ldots, h_i$ could explain e; e has some explanation; hypotheses $h_1, \ldots, h_{i-1}$ are false; *therefore* hypothesis h_i is true.

In this way eliminative induction may be seen as a limiting case of IBE, the limiting case where IBE leads to knowledge. Of course, the first premise (and indeed the second and third) will need to be discovered by the methods of scientific investigation which will be broadly inductive in the sense being used here. They may themselves be instances of eliminative induction or may depend on enumerative induction. And so the casting of eliminative induction is not intended to solve the problem of induction, but rather to reveal the structure of an important route to inductive knowledge. Ultimately any attempt to show that our inductive practices, whatever they are, can lead to knowledge will have to appeal to externalist epistemology in some form.

References

Achinstein, P. 1983. *The Nature of Explanation*. Oxford: Oxford University Press.

Armstrong, D. M. 1973. *Belief, Truth, and Knowledge*. Cambridge: Cambridge University Press.

Armstrong, D. M. 1983. *What is a Law of Nature?* Cambridge: Cambridge University Press.

Bird, A. 2005. Abductive knowledge and Holmesian inference. In T. S. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology*, pp. 1–31. Oxford: Oxford University Press.

Braithwaite, R. B. 1953. *Scientific Explanation*. Cambridge: Cambridge University Press.

Earman, J. 1992. *Bayes or Bust?* Cambridge MA: Bradford.

Foster, J. 1983. Induction, explanation, and natural necessity. *Proceedings of the Aristotelian Society 83*: 87–101.

Goodman, N. 1954. *Fact, Fiction, and Forecast*. London: Athlone Press.

Kitcher, P. 1993. *The Advancement of Science*. New York: Oxford University Press.

Lipton, P. 2004. *Inference to the Best Explanation* (2nd ed.). London: Routledge.

Mellor, D. H. 1991. The warrant of induction. In *Matters of Metaphysics*. Cambridge: Cambridge University Press.

Nozick, R. 1981. *Philosophical Explanations*. Cambridge MA: Harvard University Press.

Papineau, D. 1993. *Introduction to Philosophical Naturalism*. Oxford: Blackwell.

Popper, K. 1959. *The Logic of Scientific Discovery*. London: Hutchinson.

Skyrms, B. 1999. *Choice and Chance: An introduction to inductive logic* (4th ed.). Belmont, CA: Wadsworth Publishing.

van Cleve, J. 1984. Reliability, justification, and the problem of induction. *Midwest Studies in Philosophy 9*: 555–67.