1. Abductive Knowledge and Holmesian Inference

Alexander Bird

1. INTRODUCTION

The usual, comparative, conception of Inference to the Best Explanation (IBE) takes it to be ampliative. In this paper I propose a conception of IBE ('Holmesian inference') that takes it to be a species of eliminative induction and hence not ampliative. This avoids several problems for comparative IBE (e.g. how could it be reliable enough to generate knowledge?). My account of Holmesian inference raises the suspicion that it could never be applied, on the grounds that scientific hypotheses are inevitably underdetermined by the evidence (i.e. are inevitably ampliative). I argue that this concern may be resisted by acknowledging, as Timothy Williamson has shown, that all knowledge is evidence. This suggests an approach to resisting scepticism different from those (e.g. the reliabilist approach) that embrace fallibilism.

2. SCEPTICISM AND EVIDENCE

There is a sceptical argument that goes like this. We like to think that we are in a world not only such that we seem to see an environment of physical objects of certain sorts, but also where such objects do indeed exist and are in many respects as they seem to us to be. Such a world we may call the 'good situation'. However, says the sceptic, our evidence for thinking that we are in the good situation is just the way things seem to us to be—our subjective sense-impressions. And, says the sceptic, our evidence would be just the same if we are in the bad situation, namely, one where an evil demon or some other deceiving device causes us to have the same set of sense-impressions. Since our total evidence in the good situation is identical to what it is in the bad situation, we cannot

know that we are in the one situation rather than in the other. Hence we cannot know that there really are the physical objects there appear to be.

As I have outlined it, the sceptic's argument rests on two premises:

(EVEQ) S's evidence in the good situation is the same as in the bad situation;

(DIFF) If S is to know p then S's evidence must be different from what it would have been in any situation where $\neg p$;

from which the sceptical conclusion:

(SCEP) If S is in the good situation, S does not know that S is in the good situation; follows immediately.

There are many ways of attempting to deal with scepticism. In the context of the current sceptical argument two strategies are apparent, corresponding to the denial of the two premises, (DIFF) and (EVEQ).

(DIFF) might be denied by arguing that differences in states of knowledge are not dependent on differences in evidence alone. They may also depend on differences in facts external to the subject's evidence. In particular, knowledge is sensitive to the nature of causal connections between the subject and the environment or the reliability of the processes by which the subject acquires his or her beliefs. These facts need not be included in the content of the subject's evidence. So two individuals may have identical sets of evidence, and identical beliefs, but the one knows something the other does not, since the two are in different environments, where one is propitious for knowing and the other is not. So, S can know that S is in the good situation since in the good situation the methods that link S's evidence and S's beliefs will be reliable, which they won't be in the bad situation. Although reliabilism is not the only instance of the approach to scepticism that denies (DIFF), I shall for convenience call this approach the 'reliabilist strategy'. Since (DIFF) is incompatible with fallibilism, as usually conceived, the term 'reliabilist strategy' may do duty for any strategy that rejects scepticism while also embracing fallibilism.[1]

To deny the sceptical conclusion, we may instead reject premise (EVEQ), as has been suggested by Timothy Williamson.[2] Quite

[1] For a discussion of what fallibilism amounts to, see Reed (2002: 143–57).

[2] Williamson, (2000*b*: 625 n. 13). Note first that Williamson does not explicitly portray the sceptic's argument as proceeding from (DIFF) as well as (EVEQ). And secondly that

independent of any discussion of scepticism, there is reason to think that our evidence is just what we know (Williamson 1997):

(E = K) S's evidence is precisely what S knows.

So, if we take that on board, (EVEQ) becomes:

(EVEQ*) S's knowledge in the good situation is the same as in the bad situation.

But clearly (EVEQ*) is just what is in dispute. The sceptic's conclusion is that we can know no more in the good situation than in the bad situation, which is very little. So (EVEQ*) is no good as a premise in that argument, and likewise (EVEQ) is no good too. The sceptic just begs the question. Of course a bad argument can have two false premises and so the two diagnoses of the error in this version of scepticism are not in direct conflict. One aim of this paper is to explore the consequences and limitations of this argument as a response to scepticism that does not rely on the rejection of (DIFF). My particular focus will be upon abductive inferences. Williamson himself (2000*a*: 174) is inclined towards acceptance of (DIFF) when discussing whether the difference between relevant and irrelevant alternatives makes trouble for the sceptic: 'it is difficult not to feel sympathy for the sceptic here. If one's evidence is insufficient for the truth of one's belief, in the sense that one could falsely believe *p* with the very same evidence, then one seems to know *p* in at best a stretched and weakened sense of ''know''.' An immediate implication of Williamson's comment is that where *p* is the conclusion of a normal ampliative argument, one whose conclusion is not entailed by its premises, then that argument cannot yield knowledge of the conclusion in the proper sense of 'knowledge'. (Why I say a 'normal' ampliative argument, rather than 'any' I shall explain shortly.)

Are there then reasons for wanting to retain some version or other of (DIFF)? First, it seems that the power of scepticism would be less easily explained if we think that all its premises are at fault than if we identify just one subtle error. Secondly, (DIFF) is related to Williamson's safety condition on knowledge, that, when cases α and β are close to one another, if S knows *p* in case α then in case β S does not falsely believe

Williamson does not intend his case against scepticism to depend upon this argument. (Rather, it seems, this argument forces the sceptic to articulate a different, phenomenal, conception of evidence, that is ultimately indefensible.)

p (Williamson 2000*a*: 128). If cases where the subjects have identical evidence are classed as close to one another, then the safety principle entails (DIFF). Thirdly, given the equation (E = K) (DIFF) comes out as trivially true as it stands (although this is not the case for a diachronic version of (DIFF), as we shall see). Fourthly, (DIFF) rules out any account of knowledge of the following form, knowledge is justified true belief plus X, where the truth condition is non-redundant.[3] This is because such an account contemplates situations just like knowing with regard to justification but without truth. But (DIFF) requires situations that are like knowing but without truth to differ with respect to evidence also. Such situations will thus differ with regard to justification also (if justification is a relation to the evidence). So there are no situations of the kind JTB + X accounts envisage. Since we want to exclude such accounts anyway, thanks to Gettier's examples, a principled condition on knowledge such as (DIFF) that does so is thereby at an advantage. Finally, many people find (DIFF) compelling for the following reason. Knowing should imply epistemic responsibility, and responsibility is reasonably taken to mean appropriate sensitivity to the evidence. But if we reject the idea that there is some evidential difference between knowing and not knowing, the difference between S in the good situation who does know and S* in the bad situation who does not, has nothing to do with epistemic responsibility and has everything to do with epistemic luck. The reliabilist strategy allows that, for all else S knew, S might have been in the bad situation, but nonetheless gets to know that she is not. Many find this counter-intuitive, and it is a major part of the force behind epistemological internalism.

By retaining some difference principle we may respect these intuitions; at the same time sceptical conclusions may be resisted by rejecting the sceptic's appeal to premises akin to (EVEQ). In accordance with the Williamson strategy, (E = K) will be assumed throughout this paper.

3. RELIABILISM AND INFERENCE TO THE BEST EXPLANATION

A further reason for wanting to rest resistance to scepticism at least partly on the Williamson strategy rather than the reliabilist strategy is

[3] I am grateful to Richard Fumerton for this point.

the concern that reliabilism alone does not have sufficient resources to account for all knowledge gained by abductive inference. It is reasonable to hope that reliabilism might be able to account for knowledge gained from enumerative (or 'Humean') induction. Such is Mellor's approach, for example.[4] The difficulty is that Humean induction describes very little of the inference that takes place in science. The larger part of scientific inference is abductive. By 'abductive' inference I shall mean an inference where a central component of that inference is the fact that the inferred (purported) facts provide a putative explanation of the evidence or some part thereof. I shall treat 'Inference to the Best Explanation' (IBE) as a synonym for 'abductive inference', treating 'Inference to the Best Explanation' less as a description than as a name for a certain class of inferences that trade on the explanatory capacities of what is inferred. What exactly is involved in IBE is one of the issues to be discussed.

The application of reliabilism to IBE is thoroughly problematic. Typical accounts of IBE are what I shall call 'comparative'. They involve comparing putative explanations of some evidence with respect to their explanatory 'goodness'. Such accounts may permit or enjoin acceptance of that putative explanation which is comparatively better than all the others (hence inference to the 'best' explanation). They may possess other features, such as the requirement that the best be clearly better than its competitors and that it meet some minimum threshold of goodness. Nonetheless, the common feature is a comparison of competing possible explanations.

In order to give a reliabilist explanation of how comparative IBE can generate knowledge, two tasks must be carried out. The first is to explain what the goodness of a explanation is. The second task is to show that goodness correlates with truth. So, in fulfilling the first task, we might identify goodness with certain virtues of explanation such as simplicity, tendency to provide explanatory unification, or capacity to provide understanding (what Peter Lipton (1991: 61) calls 'loveliness')—we infer the most virtuous of the potential explanations. (A potential, or putative, explanation is, very roughly, something that would be the actual explanation if it were true.) Then, for such accounts to be plausible as explanations of inductive knowledge, the second task

[4] For a reliabilist account of Humean induction see Mellor (1987), who explicitly excludes Inference to the Best Explanation from his considerations.

requires showing that it is at least plausible that explanatory virtue is a reliable indicator of truth.

It is this second requirement that ought to be particularly worrying for the reliabilist. Many debates surrounding IBE question whether accounts of explanatory goodness make that goodness too subjective for it to be even possible for them to be correlated with truth. But even if we can show that simplicity, unification, loveliness, and the like are objective, that only shows that they might be correlated with truth, not that they are. For reliabilism to be a plausible account of knowledge via IBE, we should seek some evidence that there actually is such a correlation.

The problem is that such evidence is thin. Good explanations are frequently falsified and often replaced by less virtuous ones. The theory of relativity is less simple than the Newtonian mechanics it replaced, while many aspects of quantum theory are distinctly lacking in virtue and might even be regarded as explanatorily vicious (renormalization, non-locality, complementarity, and so on). The ancient theory of four elements was replaced by one with over one hundred elements. Even if the balance seemed to be restored by the discovery of the three subatomic components of atoms, it was put out of kilter by the subsequent discovery of a zoo of such particles. The Pessimistic Induction is overstated; nevertheless, it is true that good explanations are frequently falsified. They are often replaced by hypotheses that were earlier considered (or would have been considered) less virtuous ones. Thus although quantum theory is a better explanation of the current evidence than classical mechanics, since the latter is falsified by current evidence, matters are reversed when we consider the old evidence, that available say in the middle to late nineteenth century. It is significant that explanatory goodness is frequently in due course falsified by the evidence.

Thus, in so far as explanatory goodness (e.g. simplicity, elegance) is independent of any specific set of evidence, we find that such goodness often decreases as theories change. Even if there is some degree of correlation between goodness and truth, that correlation is, I fear, too weak to reach a level of reliability required to generate knowledge. There are differences among reliabilists about what degree of reliability is required. Some urge that if the level of reliability is less than 100 per cent, then any such account of knowledge is liable to fall foul of Gettier-style cases. Clearly inference to the most virtuous explanation does not

achieve that level of reliability. It is also doubtful whether it meets even any lower threshold that would nonetheless be a plausible degree of reliability in a reliabilist account of knowledge. It is worth recalling that the comparison ought not be simply between pairs of competing hypotheses considered individually; if comparative IBE were to be reliable, the preferred hypothesis would also have to be more plausible than the disjunction of the remaining hypotheses.

4. DIRECT AND INDIRECT EVIDENCE

The following propositions:

(a) abductive inference is comparative IBE;
(b) abductive inference can be knowledge generating;
(c) the difference principle, (DIFF), is true;

are in tension. Although they are strictly consistent, they can be jointly true only in virtue of peculiar and unrepresentative inferences. As applied to the large majority of instances of IBE they cannot be jointly satisfied.

In a comparative IBE the various hypotheses under consideration are consistent with the evidence. Considered individually in relation to the evidence, each could be true. That is why the inference needs to take into consideration the relative goodness of each explanation (whether it be simplicity, explanatory power, etc.). Comparative IBE is ampliative—the evidence does not entail the conclusion. That a knowledge-generating inference is ampliative does not immediately entail that it doesn't satisfy (DIFF). Consider an inference made by S to a proposition such as 'S exists' or 'S has some evidence' (which we shall take to be true). (DIFF) asks us to consider S's evidence in a situation where the inferred proposition is false. Clearly S's evidence would be different (namely, none at all) in such situations. Hence any inference, ampliative or not, to such propositions will satisfy (DIFF). Hence the mere fact that comparative IBE is ampliative is insufficient to show that (a)–(c) are inconsistent. However, it is clear that the cases that allow ampliative inferences to satisfy (DIFF) are unusual and the conclusion propositions in question are not the sort that one would normally employ an IBE to ascertain. Consequently, the vast majority of actual IBEs, if they are comparative, and if they are knowledge-generating,

will not satisfy the difference principle. Correspondingly, and this is the conclusion that I will be focusing on later: if IBE is knowledge-generating, and if the difference principle is to be respected, then IBE cannot be comparative. Some other account of what is going on in (knowledge-generating) IBE must be found.

At first sight a strategy for combining comparative IBE with respect for a difference principle might be the following. The facts we normally think of as evidence for a hypothesis include the results of experiments and observations, previous theoretical conclusions, and so forth. Let us call this 'direct' evidence. In a case of comparative IBE, the direct evidence relevant to a set of hypotheses will be, principally, the facts explained by those hypotheses. It is tempting to think that direct evidence exhausts the relevant evidence, but on reflection it does not. There is also *indirect* evidence. For example, the simplicity or elegance of a hypothesis might be further evidence in its favour. In general, when employing IBE, facts such as the fact that one hypothesis is a better explanation of the (direct) evidence than its competitors can be known to the investigator and hence can be part of the investigator's total evidence.

Noting that indirect evidence exists is one way of repelling the claims of the underdetermination of theory by evidence. For if hypotheses differ, for example, in their simplicity, then there will be an evidential difference between those hypotheses. One might hope to apply this to the current problem of reconciling IBE and some difference principle as follows. If we add the indirect evidence to the direct evidence, then perhaps the total evidence the subject has for hypothesis h might be evidence that S could not have in any situation where h is false. This requires that the following should hold:

$$<\text{S has evidence } e> \text{ entails } <h>$$

where e is the total evidence, indirect evidence included. But even those who are alive to the importance of indirect evidence tend still regard IBE as ampliative on this total evidence. It seems that one could always have the same total evidence and yet be mistaken in the conclusion that IBE presses upon us.

Of course the total evidence e might *determine* a single conclusion in the following sense. The total evidence, including the indirect evidence, can yield an unambiguous conclusion when IBE is applied to that evidence. It can be that a single conclusion is consistent with possession of the evidence and with the application of IBE. That means:

<S has evidence *e* & IBE is truth-preserving> entails <*h*>

can hold. Does this show how a difference principle may be respected? No it does not, unless the fact that IBE is truth-preserving is amongst S's evidence—difference principles concern themselves only with differences of evidence. If IBE is indeed truth-preserving it may be possible to know that it is. Yet we would not want such knowledge to be a condition of IBE's producing knowledge. For then we really would fall foul of the circularity charges that Hume urged.

5. THE CHALLENGE OF ABDUCTIVE INFERENCE

In §2 I presented a sceptical argument, proceeding from two premises (EVEQ) and (DIFF), and two strategies for resisting scepticism, corresponding to the denials of the two premises. The rejection of (DIFF) goes hand in hand with a reliabilist approach to knowledge. §3 argued that reliabilism is not a satisfactory way of explaining how IBE yields knowledge. Hence a resistance to scepticism as regards IBE ought to consider rejecting (EVEQ) and retaining (DIFF). In §2 I noted other reasons for wanting to retain (DIFF). However, §4 shows that if the difference principle is to be retained then at least those instances of IBE that are knowledge-yielding had better not be comparative. In the next section I shall present an account of knowledge-yielding IBE that is not comparative. Here I shall consider in general terms the strategy of rejecting (EVEQ) and retaining (DIFF), as applied to IBE.

Can we extend Williamson's anti-sceptical argument to abductive scepticism by rejecting (EVEQ)? There are obstacles to so doing. Consider an inference to the conclusion *h*. The most straightforward attempted application of the Williamson strategy would amount to claiming, ''The sceptic's premise (EVEQ) is false. It is false because in the good situation S will know *h* and so have different evidence (since E = K) from S* in the bad situation where S* does not know *h*.'' As a response to the sceptic this seems sadly inadequate. Scientific inferences of the sort we are interested in are those that are supposed to take us from a state of possessing evidence along with ignorance regarding some proposition to a state of knowledge concerning it. And so a response to the sceptic that compares the states of evidence *after* the inference seems to have missed the point. After all, if in the good situation the subject

does come to know h as a result of an inference, as the anti-sceptic maintains, h may perhaps then become evidence that can be used in favour of some other proposition. But it is no part of S's evidence for h itself. The inductive sceptic intends to compare S's evidence in the good situation with S's evidence in the bad situation *before* the inference. The claim is that it is possible for S's evidence to be the same in both. So what we need is a diachronic version of the sceptical argument, which will now proceed as follows.

A subject S with evidence e at t_0 reasons on the basis of e to a conclusion p at t_1 later than t_0. The good situation is one where e is true and p is true and the bad situation is one where e is true and p is false.

$(\text{EVEQ})_\text{d}$ S's evidence in the good situation at t_0 is the same as in the bad situation at t_0;

$(\text{DIFF})_\text{d}$ If S is to come know at t_1 that p, by inference from evidence possessed by S at t_0, then S's evidence must be different at t_0 from what it would have been in any situation where $\neg p$.

Now consider a case where S, who has evidence e, is considering rival and mutually inconsistent scientific hypotheses, p and q. Let is be that e entails neither p nor q, and furthermore, both p and q are consistent with S possessing evidence e. By $(\text{DIFF})_\text{d}$ S does not know that p at t_1.

Now the assumption of the diachronic version of the evidential equivalence claim, namely $(\text{EVEQ})_\text{d}$, is no longer question-begging. We apply (E = K) and then $(\text{EVEQ})_\text{d}$ tells us that S's *knowledge* in the good situation at t_0 is the same as in the bad situation at t_0. But that does not beg the question as to whether S comes to know p at the later time t_1 after the inference.

This would appear to cast doubt on the strategy for combating scepticism derived from Williamson.[5] Nonetheless, I do think that an

[5] For the response given above can be employed for any proposition that comes to be known as a result of an inference. And for most ampliative inferences it will appears that there can be a bad situation that makes $(\text{EVEQ})_\text{d}$ true. On the other hand, if the strategy is applied to propositions that are not inferred, but are believed directly (e.g. perceptual propositions), then it is not clear that the argument presented accurately characterizes the sceptic's position. For if the proposition is not inferred, an argument based on the nature and sufficiency of evidence seems an inappropriate way of spelling out the sceptical worry.

adaptation of Williamson's argument does play a valuable role, as I shall show later. Furthermore, the version of the sceptic's argument that it engenders, as we have just seen, prompts us to look for an account of IBE that permits the truth of the difference principle and the falsity of the evidential equality claim.

6. HOLMESIAN INFERENCE

We are apt to classify the inferences ascribed to Sherlock Holmes (such as identifying a criminal on the basis of the mud on a man's boot, the analysis of a cigar ash, and so on) as inductive. Yet Conan Doyle described Holmes's method as *deductive*. This appears to be a solecism.[6] On the other hand, Sir Arthur goes on to provide details of the method which allow for a reconciliation of this terminology. On several occasions Holmes tells Watson, ''Eliminate the impossible, and whatever remains, however improbable, must be the truth.''[7] That clearly is deductive. If Holmes starts by knowing that one of ten hypotheses is true and by dint of further evidence gathered in the course of his investigations comes to know of nine of them that each is false, the deduction tells him that the tenth must be true.

Holmes's method is deductive if and only two conditions are met:

(*a*) Holmes knows that one of the ten hypotheses is true;
(*b*) Holmes obtains evidence that is inconsistent with nine of the hypotheses.

As a procedure, what I shall call Holmesian inference has the following form. From initial evidence, e_i, Holmes gets to know that one of hypotheses $h_1, \ldots, h_n$ can be true. These hypotheses are explanatory hypotheses; they explain some subset e_s of the evidence. Holmes then collects additional evidence, e_a, such that e_a (given e_i) rules out $h_1, \ldots, h_{n-1}$. Hence Holmes may deduce that h_n is true. Holmesian inference requires three premises:

[6] This criticism is made e.g. by Lipton (2001).

[7] See, e.g. Conan Doyle A. 1953*b*: 94, 118; 1953*a*: 1089. Kitcher and Earman make favourable references to Holmes in support of eliminative induction (Kitcher 1993: 239; Earman 1992).

(i) the fact e_s has an explanation (*Determinism*);[8]
(ii) $h_1, \ldots, h_n$ are the only hypotheses that could explain e_s (*Selection*);
(iii) $h_1, \ldots, h_{n-1}$ have been falsified by the evidence (*Falsification*).

We need not see this solely in procedural terms. For we may say that at the end of the investigation Holmes's final evidence, $e_f = e_i + e_a$, entails the one hypothesis, h_n. In summary, Holmesian inference may be explained thus:

(HOLMES) S knows h by Holmesian inference from evidence e iff S deduces h from e, which includes the proposition s, where, for some $e_s \subset e$, s is the proposition that there is some explanation of e_s.

Holmesian inference is Inference to the Best Explanation. But it does not involve the selection of potential explanations according to their explanatory virtues. Knowledge by Holmesian inference is gained only when the evidence rules out all but one of the potential explanations—the best explanation is the *only* explanation of the evidence, or some part of the evidence, that is consistent with the evidence.

Now consider S who infers that p by Holmesian inference from what S knows. Since p is entailed by what S knows, it could not have been that in some other world S has the same evidence but p is false. Hence if there is a bad situation in which S infers p by Holmesian inference but p is false, S must have different evidence in that situation. Therefore (EVEQ) and $(\text{EVEQ})_d$ will not be true in this case and so this approach is consistent with Williamson's.

In what follows I shall argue that knowledge by Holmesian inference is possible. We may thereby have a view of abductive knowledge which coheres with the approach to scepticism that grants the sceptic a version of the difference principle.

The possibility of knowledge from Holmesian inference is of course controversial. It requires not only the possibility of the truth of each of the three premises *Determinism*, *Selection*, and *Falsification* but also possibility of (concurrent) knowledge of the three premises. I'll comment on each in turn. The premise *Falsification* ought to be the least

[8] I take the names of the first two assumptions, Determinism, and Selection, from von Wright 1951: 131.

controversial. On some occasions at least we are able to falsify hypotheses. If so we should sometimes be able to falsify all but one of the hypotheses in a finite set of mutually inconsistent hypotheses that includes the true hypothesis. Although the least controversial aspect of Holmesian inference, *Falsification* is not entirely uncontroversial. The (Duhem–)Quine thesis alleges that we always have the choice of avoiding a falsification by changing an auxiliary hypothesis. Knowledge from Holmesian inference requires that we know that we have falsified alternative hypotheses. Hence the Quine thesis presents a challenge to Holmesian inference if the thesis is taken to be the claim that we cannot *know* that any hypothesis is falsified by the evidence. Let us say that we do know that some observational proposition *o* and some relevant auxiliary hypothesis *a* are true, such that from *o* and *a* the falsity of the target hypothesis *h* is deducible from *o*&*a*.[9] Then the subject can know the falsity of *h*. Hence the denial that we can know that *h* has been falsified requires that we deny that the auxiliary hypothesis *a* is known (assuming knowledge of the observational proposition *o*). Consequently, the Quine thesis, if it is to be a challenge to Holmesian inference, must be regarded as stating that auxiliary hypothesis cannot be known to be true. Hence the thesis that we cannot know that a hypothesis has been falsified thus implies a general scepticism about the possibility of knowledge of (auxiliary) hypotheses. To the extent that the Quine thesis is used to question the possibility of knowledge from Holmesian inference, it begs the question by assuming a scepticism at least as strong as that which it seeks to establish.[10]

The premise *Determinism* should not be too controversial either. We need it because ruling out all but one of the potential explanations is not quite sufficient for abductive knowledge by Holmesian inference. The evidence may not require an explanation at all. So Holmesian inference

[9] Assuming E = K, we can say, as (HOLMES) requires, that the falsity of *h* is deducible from the subject's evidence. It is clearly little consolation to the supporter of the Quine thesis to argue that this is not falsification by evidence, by denying that E = K. Falsification by a known proposition is just as bad. This is why maintaining the Quine thesis requires denying *knowledge* of the auxiliary hypothesis. Parenthetically, I suggest that the Quine thesis is implicitly operating with a very restricted notion of evidence, for example, the phenomenal conception of evidence discussed in Williamson (2000*b*).

[10] Note that Susan Vineberg (1996) criticizes Kitcher's eliminativism (1993) on the ground that for him acceptable auxiliaries are determined by prior scientific practice. Here it is not merely prior practice that delivers the acceptable auxiliaries but rather that the relevant auxiliaries are known.

requires that the subject know that there is *some* explanation or other. The subject must have some evidence that rules out the null hypothesis, that there is no explanation. In some cases the null hypothesis may indeed be true. We may take one lesson of quantum indeterminacy to be just that. But in many cases the existence of some explanation is not in doubt. The detective knows that the bullet didn't just materialize out of nothing in the brain of the victim and that the entry wound didn't just come from nowhere. Similarly, that there is some correct explanation or other for the extinction of the dinosaurs is not up for question even if the nature of that explanation is. The assumption of *Determinism* is not the false assumption that universal determinism holds, but the assumption that some part of the evidence in question has an explanation, which may be true for many cases. Van Fraassen (1980: 21) says that ''the realist will need his special extra premise that every universal regularity in nature needs an explanation, before the rule [of Inference to the Best Explanation] will make realists of us all.'' But such a strong, universal premise is not required. It is sufficient for Holmesian inference that we know of *some* facts of interest that they have an explanation.

The premise *Selection* of the Holmesian inference is the most controversial. Most philosophers of science are not willing to grant that the hypotheses that could explain some piece of evidence may be finite in number. They accept the thesis that theories are radically underdetermined by the evidence. I tackle this problem in §§ 8–11. Before that I shall compare Holmesian inference with another, related version of eliminative induction.

7. PAPINEAU ON NON-ENUMERATIVE INDUCTION

David Papineau (1993: §5.15) proposes a similar model of induction, based on Mill's methods.[11] I have characterized Holmesian inference as employing three kinds of premise: *Determinism* (the fact e_s has an explanation), *Selection* ($h_1, \ldots, h_n$ are the only hypotheses that could explain e_s), and *Falsification* ($h_1, \ldots, h_{n-1}$ have been falsified by the evidence). Papineau regards *Falsification* as the only premise in an argument that leads to the conclusion that some h_n is true. According

[11] Von Wright's account of eliminative induction (1951) also starts from a consideration of Mill.

to Papineau and von Wright *Determinism* and *Selection* are not premises of the subject's reasoning at all. This is why, on their view, inductive reasoning of this kind is ampliative and not deductive. Instead, according to Papineau, the subject is simply disposed to assert h_n once the subject knows $h_1, \ldots, h_{n-1}$ to be false. If *Determinism* and *Selection* are true (even if not known to be), then that disposition will be reliable. So overall, given knowledge of *Falsification*, the subject's belief in h_n will be reliably formed and hence will be knowledge.

Papineau's account is satisfactory only if we can regard the disposition in question as part of the process or rule whereby h_n was inferred *rather* than as masking undischarged premises of the form of (i) and (ii). This is a problem that requires a general answer. Consider a subject who reasons as follows: premises P, $P \rightarrow Q$, conclusion Q. If this subject knows P but does not know $P \rightarrow Q$ (even though $P \rightarrow Q$ is true), then we must deny knowledge to this subject of the conclusion Q. Now consider a second subject, who also does not know $P \rightarrow Q$ and who argues from the single premise P to the conclusion Q, being disposed that way. In the latter case, if Papineau's view is correct, the subject will get to know that Q. But is the second subject really entitled to the status of knowledge of Q that is denied to the first? Furthermore, the two cases are not so clearly distinct, since someone who believes that $P \rightarrow Q$ will be disposed to believe Q when they believe P (indeed on some account that disposition is partly constitutive of belief that $P \rightarrow Q$). So the case of believing $P \rightarrow Q$ will include the case of being disposed to infer Q from P. If the latter gives knowledge one might imagine that the former should also.

A further problem for Papineau concerns the nature of the disposition that has to exist in lieu of the premise (ii). We have many innate cognitive dispositions—these are typically fairly general in nature and can be explained by their evolutionary contribution to fitness. We acquire further cognitive dispositions through experience of the world. Typically these dispositions will be cognitive habits, acquired by repeated use or by repeated experience. In the sort of case we are considering neither of these apply. One of Papineau's examples is the identification of the human immuno-deficiency virus as the agent that causes AIDS. This he presents as being achieved by the elimination of other possible viruses as candidates for the agent. So the subject who thereby gets to know that HIV is the cause of AIDS must be disposed to infer that HIV causes AIDS when that subject knows that other viral

candidates have been eliminated. Such a disposition is clearly one that could not be innate. Could it be acquired by repeated employment or by experience of constant conjunction? Surely not—the disposition is far too specific for that. And in any case this is a disposition that must be able to exist before the subject makes first inference of the relevant form. So where does the disposition come from? Presumably it comes from background beliefs the subject has concerning microbiology. It will be generated by beliefs such as the beliefs that only viral infections do not respond to antibiotics and that AIDS does not respond to antibiotics. It is difficult to see how a very specific disposition linking beliefs of a theoretical kind that is brought about by beliefs with theoretical content is itself vary far short of being a belief, even if only a tacit belief. And if this is the case that belief or quasi-belief will look much closer to an undischarged premise than to a mere disposition, part of the process and not part of the content.

Why does Papineau want to avoid the suggestion that (i) and (ii) might be genuine premises? Presumably, I surmise, because he thinks that there are problems concerning knowledge of (i) and (ii). One ground for doubting knowledge of (ii) I address below—this is the concern that there might be too many competing potential explanations for it to be possible to know some premise that states that they are all the potential explanations there are. Papineau does not raise this problem. Presumably it would difficult to have a disposition whose nature covered a vast range of hypotheses. Both Papineau and I require the range of hypotheses at stake to be manageable. Papineau's worry is instead a different one. Knowing that the cause of AIDS is one of viruses $v_1 \ldots v_n$ requires *knowing* that only viral diseases do not respond to antibiotics. Papineau does not deny that this could be known. But it will be known as the result of another eliminative inference of the same kind, one that asks which agents are responsible for antibiotic resistant infections. This seems to threaten some kind of regress, which is avoided by not requiring the subject to have (ii) as a premise. The philosopher can show that the disposition to assert h_n in response to knowing (iii) is a reliable one, in a naturalistic fashion, by citing the relevant facts concerning which agents are responsible for antibiotic resistant infections.

However, it is not clear that the regress in question is a vicious one. It seems that in many cases the relevant premises of the form of (ii) are known to the investigator, and in sophisticated cases, such as those just considered, they need to be. The regress is contained by noting

that for some Holmesian inferences the relevant premise (ii) need not be gained by an antecedent Holmesian inference. Later I shall consider some everyday cases where the limited range of hypotheses is knowable by simple common sense. The core of common sense may be innate knowledge or reliable intuition, which in this case constrains the range of possibilities we consider. It is clear that Holmesian inference must be supplemented by another source of knowledge of general propositions.[12]

8. THE UNDERDETERMINATION CHALLENGE TO KNOWLEDGE BY HOLMESIAN INFERENCE

Traditional approaches to inductive scepticism have ruled out anything like inductive knowledge by Holmesian inference. The common view is that inductive inference, including abductive inference, is ampliative. By definition, the conclusions of ampliative inferences are not entailed by the evidence from which they are inferred. Holmesian inference is not ampliative—the conclusion may be deduced from the three premises, *Determinism*, *Selection*, and *Falsification*. We have considered *Determinism* and *Falsification*. The claim that abductive reasoning is always ampliative typically rests on a thesis that rejects *Selection*, namely the thesis that hypotheses are underdetermined by the data. The precise nature of that thesis is itself debatable; the version I shall consider is:

(UD) There is always more than one explanatory hypothesis consistent with the evidence.

This thesis entails the claim that inductive inferences are ampliative. It is clear that knowledge by Holmesian inference is inconsistent with (UD). What reason is there to believe (UD)?

There are two considerations or kinds of consideration that are typically cited in favour of (UD). The first states that there must be such a *quantity* of distinct possible causal histories that however much evidence is gathered that rules out some of these, there will always remain more than one. The second consideration is that there be a *qualitative*

[12] One would like to show that such sources also satisfy the difference principle—or that it does not apply, e.g. by showing that the relevant knowledge is delivered by a quasi-perceptual faculty, as I suggest below.

difference between the evidence and the facts constituting the possible causal histories. If the number of possible disjoint causal histories were not too great, (UD) might still be true if the only possible evidence were of such a *kind* that it could not rule out any of these histories. I shall look at the qualitative thesis before returning to the quantitative thesis. In both cases I shall argue that the underdetermination thesis is false.[13]

9. THE QUALITATIVE THESIS OF UNDERDETERMINATION

Evidence will qualitatively underdetermine theory if evidence propositions are all of one kind and theoretical propositions are all of another kind, such that the latter are epistemically inaccessible from the former. For example, if our evidence propositions consisted solely of propositions from pure mathematics and the theoretical propositions in question concern organic chemistry, then one would not expect to be able to get knowledge of the latter by any form of inference from the former. Empiricism in one of its guises holds that our evidence propositions are always observational. A sceptical conclusion concerning the knowability of propositions concerning the unobserved may be drawn, employing the following argument:

(OBS) All evidence is observational;
(INF) From observational premises only observational conclusions may be rationally inferred;[14]
therefore
(SCEP) Only observational propositions can be known.

Many of the failings of empiricism have been adequately addressed elsewhere. Here I shall add to those arguments one that derives

[13] Other supporters of eliminative induction accept the underdetermination thesis but hold that we are able to pare down the infinite range of logically possible hypotheses to a manageably finite number. We have seen Papineau's appeal to a disposition that fulfils this function. For Kitcher (1993: 248) it is prior scientific practice that performs this function.

[14] To my mind this premise is itself highly questionable. It would be a little less questionable if we were to replace 'inferred' by 'deduced'. If we do so, then the sceptical conclusion follows only if we add a further premise to the effect that only deductive inferences can lead to knowledge (which would be acceptable to a supporter of (DIFF)). However, my strategy here is not to take issue with (INF) but with (OBS), for which reason I am happy for (INF) to be as strong as any sceptic could wish for.

from Williamson's strategy as outlined at the opening of this paper. We apply Williamson's equation of evidence with knowledge, (E = K), to the premise (OBS). This yields:

(OBS*) All knowledge is observational;

which is equivalent to the sceptical conclusion. Hence the very limitation of evidence to observational propositions is to assume what the sceptic is seeking to prove. The sceptical argument of qualitative underdetermination is question-begging.

This strategy generalizes to any argument of an analogous nature that appeals to an underdetermination of theory by evidence on the grounds of a difference in kind between evidence propositions and theory propositions. A qualitative underdetermination argument may have the following form:

(EV) All evidence is of kind *K*;
(INF) From premises of kind *K* only conclusions of kind *K* may be rationally inferred;
therefore
(SCEP) Propositions of a kind other than *K* cannot be known.

Applying (E = K) to (EV) gives (SCEP) immediately. Any such argument will be question-begging.

10. THE QUANTITATIVE THESIS OF UNDERDETERMINATION

This consideration in support of (UD) is that the range of possible explanations is too large for any possible gathering of evidence to pare that range down to one. While theory actually is often underdetermined thanks to insufficient evidence, it needs substantial argument to show that it always must be. While the qualitative consideration in favour of (UD) employed some principled (but flawed) arguments, this quantitative consideration is supported less by positive argument than by a sense, bred by familiarity with sceptical scenarios, that however much evidence one adduces in favour of an hypothesis, one could always imagine some competing hypothesis consistent with the same evidence.

A time-slice through the causal history of an explanandum constitutes in Hempel's terms a complete (rather than an elliptical) explanation.

Holmesian inquiry need not be expected to determine one complete explanation, but may hope to show that some possible fact of interest (a hitherto unknown, typically explanatory fact) is contained in every complete potential explanation that is consistent with the evidence. (UD) is true and rules this out only if, whatever one's evidence, for any fact there is always some complete potential explanation that does not include that fact. Let us suppose in accordance with the conclusions of the last section that there is no restriction on the kind of evidence available. And let us suppose, reasonably enough, that there is no finite upper limit on the quantity of evidence we may collect. (UD) then requires that for any possible fact of interest, F, there be an infinite number of distinct complete potential explanations of some explanandum not containing F, all of which are consistent with the evidence.

This is a strong requirement that we tend in everyday circumstances to think is false. A simple case is that where we know that some fact, F, exists but want to know whether it is part of the *causal* history of some explanandum, E. Mill's method of difference tells us to consider a parallel case, the foil, that has the same total history as E, except for the absence of F. If in the parallel case there is no parallel to E itself, we may deduce that either E has no explanation or that F is part of its causal history. Hence we may know the latter, if true, given that we know the premise *Determinism* (that E has some explanation). The foil, in effect, is a way of excluding all potential explanations that do not include F.

In other cases we may not know whether the possible fact in question occurred at all, and so Mill's method is not applicable. Even so, there are certainly occasions when we naturally think that one of only a finite number of potential explanations must be true. Detective stories of a kind less sophisticated than Conan Doyle's trade on this fact. Often they will involve a murder in an isolated country house or inaccessible island, where there are only so many guests, butlers, and detectives. The number of potential murderers is finite and even if we consider the possibility of more than one murderer, there is still only a finite number of mutually exclusive hypotheses concerning the identities of those responsible. Consider also the following more mundane example. I pour milk into a tall glass and leave it on the kitchen table. I leave the kitchen for a few moments and then hear a crash. I return to the kitchen to see a broken glass on the floor with milk spilt on the floor and table. The cat is standing on the table licking at the pool of milk. Let us now look for the explanation of the spilt milk. The obvious potential

explanation is that the cat knocked over the glass, which rolled off the table onto the floor. Holmesian inference says that we know that this is what happened only if there is no other explanation consistent with the evidence. Is there no other explanation here? Perhaps some other large object hit the glass and knocked it over—perhaps the dog or a cookery book falling from a shelf. But the dog is outside. So that hypothesis is falsified. All heavy objects like cookery books are in their place. Nothing like that is found on the table or floor. Furthermore the table is in the centre of the kitchen nowhere near the path of a falling object. The ceiling, by the way, is intact too. So that class of hypotheses is falsified as well. Perhaps something shook the table. Might it have been the cat? No, since the table is a heavy oak table, too heavy for the cat to move or judder. Perhaps an earthquake? But I know we don't have powerful earthquakes in South-West England, and even if we did, I would have felt one powerful enough to shake the table and knock the glass over, which I did not.

Just as in the typical country house murder, the simple example just given strongly suggests that what I know can rule out all explanations bar one. And for that reason, I know it was the cat that knocked the glass over. In contrast, consider for a moment an additional surmise, that the cat knocked the glass over as a result of trying to get the milk inside. It seems a pretty good explanation. There are other explanations, that the cat knocked the glass with its tail or just by sitting down on the table too close to the glass. These are not such good explanations. It may be that what I know makes such explanations unlikely, and may even justify my surmise. But, I suggest, if I know nothing that rules out these other explanations, then my surmise will not amount to knowledge.

11. UNDERDETERMINATION—SCEPTICAL SCENARIOS

An objector might try more abstruse explanations of the broken glass and spilt milk. Perhaps an evil demon is playing a trick on me; the cat is innocent and the demon pushed the glass over. Does anything I know rule that hypothesis out? Yes, what I know may very well rule out that hypothesis. Most obviously I might know that evil demons of that sort do not exist. Throughout science there is no evidence that such things exist. And, given the exhaustive nature of science, there would be such evidence if they did exist. Since there isn't, we know they don't. Perhaps

some might regard this as too optimistic; perhaps there is a demon who has decided to wait until precisely this moment to engage in interference with the world.

There may be enough evidence to rule out this demon hypothesis. Let us assume I know some basic physics—folk physics may be enough. If I know it, then, thanks to (E = K), it is part of my evidence. To knock a glass of milk over requires the transference of a certain minimum amount of energy to the glass. To generate that energy there must be a force acting over a certain distance. The dimensions of the kitchen and the fact that the doors and window are closed put an upper limit on the distance and so a lower limit on the force. But I know that there is nothing in the kitchen to generate a sufficient force. For instance, the energy could have been transferred by a largish object (the size of a cat) moving slowly. But no large object other than the cat was found in the kitchen. Perhaps the evil demon and its tools are invisible. But that too would require a violation of reasonably basic physical truths which I know. Similarly, momentum might have been transferred to the glass by a small object moving at greater speed (a squash ball, for instance, which has escaped my notice). But there is nothing to have accelerated the ball sufficiently. Such considerations are a convoluted way of illustrating the point that more bizarre explanations and sceptical alternatives just as much as the plausible potential explanations are ruled out by facts I know.

Furthermore, the existence of such a demon, even if not active, may be inconsistent with other knowledge that constitutes my evidence. On one view, for instance, knowledge requires reliability of nomic rather than statistical strength. The demon's existence therefore would render the mechanism of perceptual belief formation unreliable and so would rule out much perceptual knowledge. Since we are allowing perceptual knowledge, we may argue by modus tollens that such a demon does not exist. That is, if my evidence includes any perceptual evidence, then the demon explanation of the glass falling over is ruled out. Of course, the critic may suggest that the reliabilist's condition on knowledge is too strong, but if conditions on knowledge are weakened, then it may be that knowledge of the demon's non-existence may be obtained directly.

The sceptically minded critic might be tempted to point out that I might not know folk physics and other things I claimed to know in the last paragraph. Indeed those claims might all be false. Perhaps what we take to be the laws of physics are mere regularities that exist only

thanks to the will of the demon and which may be violated by the demon at will. Such an objection misses the point. Certainly, we might live in a world in which the things I have cited as evidence are not known and so are not evidence. In such a world I will not be able to get to know that the cat spilled the milk. But there are also worlds in which I do know these things, and so these things are evidence. The challenge we are currently considering is whether the sorts of thing we normally count as evidence, if they are evidence, could ever be enough to rule out all explanatory hypotheses but one. The claim I have argued for is the conditional one: *if* what we normally take to be evidence is evidence, *then* we can gain knowledge by Holmesian inference. A form of scepticism that argues that what we normally take to be evidence is not in fact known, does not undermine this conditional claim. A die-hard sceptic will regard the antecedent as excessively strong. But that reaction just adds to the plausibility of the conditional. Precisely because the antecedent is inconsistent with sceptical hypotheses, those hypotheses are ruled out by the assumption of the truth of the antecedent. Our 'normal' evidence not only rules out ordinary hypotheses; rather, the very possibility of normal evidence is incompatible with sceptical hypotheses also. Consequently those sceptical hypotheses do not support (UD) and do not undermine the possibility of Holmesian inference.

12. ATTENUATED VERSIONS OF HOLMESIAN INFERENCE

Even so, there may be some residual concern that without going as far as considering demon-laden sceptical scenarios there may nonetheless remain abstruse and unusual explanatory hypotheses that the investigator has not considered and which have not been ruled out by direct refutation. While I am not sure that this must always be the case, it is worth mentioning a possible response. This draws upon Peter Lipton's account of IBE (1991: 61). According to Lipton the investigation and ranking of hypotheses takes place only concerning live, plausible options. He thinks that there are indeed many other potential explanations out there, but these never get consciously considered. This does not matter; these are explanations that had they been considered would have been given a very low ranking. It might be that an intelligent and experienced investigator is reliably disposed to ignore only poor

potential explanations. Any potential explanation that would be a reasonably good explanation does get considered. To Lipton's picture we could add a reliabilist coda. It might be that IBE as traditionally considered, as a ranking of hypotheses according to virtue, is reliable at least as far as excluding bad explanations. I argued that we might have reason to think that traditional IBE is not sufficiently reliable to give us knowledge of its favoured hypotheses. But it may be reliable enough at giving us knowledge that very bad explanations are false. If that is the case, an investigator who ignored such explanations might not have his reliability impugned by that fact, so long as his ignoring them is reliably related to their being very bad explanations.

This attenuated version of Holmesian inference can account for knowledge by reliabilist criteria. But does it not give up on $(\text{DIFF})_d$ in the process? I am not sure that it does. As Lipton sees it, IBE is a two-step process. The first step in the process is the one of thinking up and selecting the plausible potential explanations and filtering out the implausible ones. The second step is that of ranking and selecting among them. (According the Holmesian inference, the second step is that of eliminating all but one.) Assuming reliability in the first step, the investigator is in a position to know, before embarking on the second step, that the actual explanation lies among the potential explanations now under active conscious consideration; he knows that all the others are false. Now the first step need not be thought of as a process of *inference* at all. Ignoring the very bad potential explanations is a quasi-intuitive skill; it is the product of experience not of ratiocination (Papineau's remarks about dispositions to believe can apply here). Hence the knowledge that the actual explanation is among those under consideration can be seen as akin to acquiring new evidence by observation. No process of ampliative inference was used. $(\text{DIFF})_d$ is thereby respected.

For this to work, it had better be that the actual explanation is indeed not among the explanations not considered. That it very frequently will be amounts to the concern that Lipton calls 'underconsideration'. I believe that Lipton's answer to that problem is right (Lipton 1993). He argues that we could not have even a reliable ordering of considered hypotheses unless our background theories, used to assist in this ordering, were true or approximately true. If the background theories can be true that shows that in their cases we *did* consider the true hypothesis.

Hence, so long as we can reliably rank our hypotheses for goodness, underconsideration cannot be endemic. In my view IBE does not work simply by ranking hypothesis, instead it works by refuting them (all but one). But the argument is the same. Refutation will depend on auxiliary hypotheses. For refutation to be possible the auxiliary hypothesis must be true or approximately so. Hence if refutation of some hypotheses is possible, it cannot be that the true hypothesis is never among those considered.

Another attenuation one might make to accommodate abstruse and implausible hypotheses that are nonetheless consistent with the subject's evidence amounts to a weakening of the difference principle. We have been working with the following difference principle:

(DIFF) If S is to know p then S's evidence must be different from what it would have been in any situation where $\neg p$.

This regards as relevant situations that may be vastly unlike the actual one. However, a plausible line of epistemological thought suggests that it is a philosophical illusion that knowledge is sensitive to distant possibilities. If we think of knowledge in terms of (the denial) of luck or in terms of safety, we are not obliged to focus on any more than nearby possibilities. In which case we can employ the weaker difference principle:

($\text{DIFF}_{\text{weak}}$) If S is to know p then S's evidence must be different from what it would have been in any nearby situation where $\neg p$.

Such a conception of knowledge renders Williamson's anti-sceptical strategy redundant—but we've already seen reason to doubt its efficacy. More importantly, even this weaker difference principle gives us reason to prefer Holmesian inference to comparative IBE. The latter is insufficiently reliable for knowledge even when we restrict our attention to nearby worlds. That unreliability is revealed by, for example, the history of science rather than consideration of abstruse possibilities. Restricting our attention to nearby possibilities will allow us, in some cases at least, to ignore abstruse hypotheses and to hope for evidence that will falsify all remaining hypotheses but one. In the next section I will say more about the relationship between Holmesian inference and comparative IBE.

13. INFERENCE TO THE BEST EXPLANATION RECONSIDERED

The Holmesian picture of scientific inference allows for a greatly simplified understanding of Inference to the Best Explanation. Above I stated that there are two puzzles for accounts of Inference to the Best Explanation that regard such abductive knowledge as resulting from selecting an otherwise radically underdetermined theory on the basis of explanatory goodness. The puzzles were first to give an account of this goodness, and secondly to demonstrate a correlation between it and truth. We can now see that Inference to the Best Explanation considered as Holmesian inference may eliminate these difficulties. In the first place we may construe goodness simply as not being falsified by the evidence when other hypotheses are; the best explanation will be the only one that could be true. Secondly, Holmesian inference guarantees truth when arguing from known premises.

The description just given concerns the circumstance where all but one of the hypotheses actually entertained has been refuted. The Holmesian deduction permits us to infer the truth of the remaining hypothesis. When that is the case goodness entails truth. However, we may frequently want to make inferences where we have not yet refuted all hypotheses but one. Take a simple case where two hypotheses remain unrefuted, h_1 and h_2. Let it also be the case that there is some proposition p entailed by h_1 and denied by h_2. It might be that we are not in a position to ascertain the truth or falsity of p; nonetheless, independent knowledge tells us that p is highly unlikely. Thus we may not be in a position to know that h_1 is false, and h_2 true; but we can know that this has a high probability. Thus the structure of Holmesian inference allows room for probabilistic reasoning concerning hypotheses. Similarly, someone might have evidence which while it does not rule out a hypothesis, justifies the belief that it is ruled out. Correspondingly one might come to a justified belief, by Holmesian inference, that a hypothesis is true, even if that belief does not amount to knowledge.

The probabilistic use of Holmesian inference may be available when propositions such as p in the above are of a kind previously known to us. But where we are dealing with novel, unobservable, or previously unobserved circumstances, the background information required to give us a reasoned assessment of the likelihood of refuting facts may

be unavailable. As discussed, proponents of Inference to Best Explanation have appealed to considerations of explanatory goodness of a kind that do not entail truth in order to assess the chances of hypothesis being true. Let us call this sort of goodness 'virtue'. Virtues includes features such as simplicity, unification, and Lipton's 'loveliness'. Our concern was that virtues correlate with truth too weakly to provide knowledge. It may be that in extreme cases, where one explanation is very much more virtuous than its competitors, we can know, by reliabilist criteria at least, that the loveliest explanation is true, in the absence of evidence refuting all competitors. In other cases we have at best only a justified epistemic preference weaker than knowledge. Recalling the distinction between direct evidence for a hypothesis, which is evidence, for example, that refutes a rival (or entails the hypothesis), and indirect evidence, which is evidence of explanatory virtue, then the relationship between Holmesian inference and inference to the most virtuous explanation (e.g. as understood by Lipton) may be characterized as follows. Clearly both sorts of evidence may be relevant. But direct evidence takes priority. Once refuted, a hypothesis is out of consideration, however explanatorily lovely. Indirect evidence is therefore relevant only amongst unrefuted hypotheses. Holmesian inference corresponds to the case where there is sufficient direct evidence that indirect evidence is not needed. These are the cases that yield knowledge of hypotheses. Inferences to the most virtuous explanation, such as Lipton's inference to the loveliest explanation, concern cases where there is insufficient direct evidence. These cases may yield a rational preference but typically not knowledge.

It is not to give in to scepticism to accept that many well-favoured hypotheses do not yet constitute knowledge. It is in the nature of scientific enquiry that theories concerning a subject are proposed well in advance of there being sufficient evidence to decide their truth. Early on it will be important to gather evidence that will assist in coming to such a decision. At the same time a fruitful and promising theory will become the basis of research into yet further hypotheses. Scientists, both as individuals and as a community, will need to decide whether their efforts should be put into confirming the basic theory or into research that takes that theory as a given. They will need to take a calculated gamble. The new research will be more exciting and provide greater opportunities for personal satisfaction and professional advancement. At the same time, there is the danger that they will be pursuing a

wild goose, should the basic theory turn out to be false. There is no reason to suppose that the point at which scientists decide to accept the gamble coincides with the point at which evidence is sufficient for the theory to be known, especially when we consider that the scientists will typically not know the point at which they come to know the theory. It seems a reasonable speculation that scientists will accept the gamble well *before* the theory becomes known. Even if all scientists in the field accept the gamble and give up the search for further confirming evidence, it is likely that the new research generates results that themselves are confirming of the theory, ruling out alternatives.

14. FALSIFICATIONISM

Holmesian inference may explain the surprising readiness of professional scientists to endorse Popperian falsificationism. That endorsement is surprising because, as so many have pointed out, falsificationism explicitly denies the possibility of inductive knowledge, at least as 'knowledge' is normally understood (as being factive, entailing truth), and implicitly entails an even greater degree of irrationalism about science than this. It cannot be that scientists are attracted by the whole package—it must be some label or slogan on the wrapping. There are two features of Popperianism that are attractive to scientists and which are shared by the model of Holmesian inference. The first is the simple idea that science proceeds by falsifying hypotheses. For scientists good and interesting evidence is not a pile of confirming instances. Rather it is evidence which might refute some live option. Only Popper's philosophy of science emphasized the falsification of theories. But Holmesian inference emphasizes falsification too, for it is only by the refutation of rival hypotheses that a given hypothesis can get to be known.

The second feature is that discussed in the preceding section. On the Holmesian model, scientific knowledge is available, but not too easily. Science may be in a position where our favoured theories are not known and so, in that sense, are tentative. It may be that Popper's scepticism struck a chord here too. Scientists are indeed reluctant to claim for their best theories the status of knowledge; instead they will say things like 'this is currently our best model'. For Popper himself tentativeness and

the lack of factive knowledge are features of *all* scientific belief. Which is surely absurd. There is no reason to think of the genetic theory of inheritance and the double-helix account of DNA as mere models. We shouldn't deny knowledge of the electro-magnetic nature of light or the atomic constitution of matter. But in fields still being explored, underdetermination may be the case and tentativeness will be the appropriate attitude to take. Furthermore, it may be added, even when evidence does rule out all but one hypothesis, scientists may not immediately be aware of the fact. That is, they may know a hypothesis is true, or be in a position to know it, while still being far from knowing that they know. Hence they may not be entitled to assert that they know (Williamson 1996).

15. CONCLUSION

The argument I have given rests in large part on accepting Williamson's equation of evidence with knowledge. Relative to some conceptions of evidence (e.g. that evidence is what one believes) this equation is restrictive. Clearly such a conception is inadequate, as Williamson shows. In any case, if it were true, there would be no reason to suppose that the underdetermination thesis true. Trivially, a unique hypothesis can easily be determined by what I believe—so long as I am happy to believe enough. At the same time, Williamson's equation is reasonably generous. For example it does not require certainty for evidence. Nor does it permit, as we have seen, a limitation of evidence to observational knowledge alone. It does permit inferred knowledge to be evidence.[15] I have thus assumed throughout that the propositions that we generally take ourselves to know can be regarded as among our evidence.

Armed with this reasonably rich stock of evidence we can tackle the alleged sceptical problem of inductive knowledge that argues that a difference condition on knowledge cannot be met thanks to radical underdetermination of theory by evidence. Denying the ubiquity of underdetermination of theory by evidence enables us to assert the

[15] Strictly, Williamson's arguments for E = K do not rule out a limitation of evidence to non-inferential knowledge. This lacuna is readily filled—see Bird (2004).

possibility of knowledge by Holmesian inference. While many favoured theories at the leading edge of science probably are currently underdetermined by data, that is a scientific not a philosophical problem. Some hypotheses are not underdetermined by evidence and so can be known by Holmesian inference while adhering to some version of the difference principle.[16]

REFERENCES

Bird, A. (2004) 'Is Evidence Non-Inferential?', *Philosophical Quarterly*, 54: 252–65.

Conan Doyle, A. (1953*a*) 'The Adventure of the Bruce-Partington Plans', *The Complete Sherlock Holmes*, ii (New York: Doubleday).

——(1953*b*) 'The Sign of Four', *The Complete Sherlock Holmes*, i (New York: Doubleday).

Earman, J. (1992) *Bayes or Bust* (Cambridge, Mass.: MIT Press).

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

Lipton, P. (1991) *Inference to the Best Explanation* (London: Routledge).

——(1993) 'Is the Best Good Enough?', *Proceedings of the Aristotelian Society*, 93: 89–104; reprinted in D. Papineau (ed.), *Philosophy of Science* (Oxford: Oxford University Press, 1996).

——(2001) 'Is Explanation a Guide to Inference? A Reply to Wesley Salmon', in G. Hon and S. Rakover (eds.), *Explanation: Theoretical Approaches and Applications* (Dordrecht: Kluwer), 93–120.

Mellor, Hugh (1987) 'The Warrant of Induction', in D. H. Mellor, *Matters of Metaphysics* (Cambridge: Cambridge University Press, 1991).

Papineau, D. (1993) *Philosophical Naturalism* (Oxford: Blackwell)

Reed, B. (2002) 'How to Think about Fallibilism', *Philosophical Studies*, 107: 143–57.

Van Fraassen, B. (1980) *The Scientific Image* (Oxford: Clarendon Press).

Vineberg, S. (1996) 'Eliminative Induction and Bayesian Confirmation Theory', *Canadian Journal of Philosophy*, 26: 257–66.

Von Wright, G. (1951) *A Treatise on Induction and Probability* (London: Routledge & Kegan Paul).

[16] I am grateful to Richard Fumerton and Timothy Williamson for helpful comments, as well as audiences at Dartmouth College, and the Universities of Nottingham, Cambridge, Reading, and Lund.

Williamson, T. (1996) 'Knowing and Asserting', *Philosophical Review*, 105: 489–523.

——(1997) 'Knowledge as Evidence', *Mind*, 106: 717–41.

——(2000*a*) *Knowledge and its Limits* (Oxford: Oxford University Press).

——(2000*b*) 'Scepticism and Evidence', *Philosophy and Phenomenological Research*, 60: 625.