E & ¬H

A significant issue arises in epistemology which I call the *problem of misleading evidence*. Suppose you have good inductive evidence E for the hypothesis H, and you are thereby justified in believing H. If (E & ¬H) obtains, then your evidence E is misleading insofar as E justifies H, but H is false nevertheless.[1] Put in the simplest terms, the problem is to account for how you can be justified in believing that your evidence isn't misleading, i.e., how you can be justified in believing that ¬(E & ¬H).[2]

My treatment of the topic will proceed as follows: §1 sets out the problem of misleading evidence and how it can serve as the basis for an argument that justification isn't closed under known logical implication. §2 defends the closure principle by affirming that we can have justification for believing that our evidence isn't misleading. §3 makes the case that such justification is empirical, not a priori. §4 deals with the relation between the problem of misleading evidence and choice between empirically adequate hypotheses; such choice doesn't (always) rest on a priori considerations, as many have thought. §5 responds to an argument, similar to Roger White's well known objection to dogmatism, that there can't be empirical justification for believing that one's evidence isn't misleading. §6 defends dogmatism against White's objection, and proposes a solution to a certain version of skepticism about induction. A theme throughout the paper is that two widely accepted theses about justification, which I call the *Entailment Principle* and the *Confirmation Principle*, are false.

1. Misleading evidence, epistemic closure, and skepticism.

The possibility that your evidence is misleading creates trouble because the evidence you have, E, is entailed by (E & ¬H). Loosely speaking, E is exactly the evidence you would expect to have if (E & ¬H) were true. So, it seems that E can't be evidence against (E & ¬H) or, equivalently, evidence for ¬(E & ¬H). To be a bit more precise: If X entails Y, X could be true only if Y were true. So, if X entails Y, Y isn't evidence that X is false. Given the assumption that Y must be evidence for ¬X for Y to justify ¬X, we have:

Entailment Principle. If X entails Y, then Y doesn't justify ¬X.

In particular, since (E & ¬H) entails E, E can't provide empirical justification for accepting ¬(E & ¬H).

The Entailment Principle is related to another thesis which is prominent in formal epistemology:

Confirmation Principle. Y justifies X only if Y *confirms* X. That is, Y justifies X only if Pr(X/Y) > Pr(X).[3]

The thought here is that if Y is to justify one in believing X, Y ought to increase how much reason one has to believe X. How much reason one has to believe X, apart from any reason provided by Y, is supposed to represented by the unconditional probability Pr(X). How much reason one has to believe X, given that Y holds, is supposed to be represented by the conditional probability Pr(X/Y). So, if Y increases the reason one has to believe X, Pr(X/Y) exceeds Pr(X). It follows that Y justifies X only if Pr(X/Y) is greater than Pr(X), as the Confirmation Principle requires. Now, it is a theorem of the probability calculus that if X entails Y, then Pr(¬X/Y) ≤ Pr(¬X).[4] This fact combined with the Confirmation Principle implies that if X entails Y, then Y doesn't justify ¬X. That is just what the Entailment Principle says. In other words, given certain commitments as to how epistemic probabilities are related to facts about justification, the Entailment Principle is a consequence of the Confirmation Principle. I will return to the role and status of the Confirmation Principle in §5.

Assume that E is your total relevant evidence. The Entailment Principle implies that E can't be evidence for ¬(E & ¬H). So, under these circumstances, you have no empirical justification for ¬(E & ¬H). In addition, whether ¬(E & ¬H) holds is a contingent matter. Many philosophers hold that there can't be a priori justification for believing contingent propositions. Call this claim *Hume's Principle*. It follows from Hume's Principle that you have no a priori justification for believing ¬(E & ¬H).[5] Putting these points together, you have neither empirical justification nor a priori justification for believing ¬(E & ¬H), so you aren't justified in believing ¬(E & ¬H) at all.

Matters worsen quickly. A widely endorsed principle about epistemic justification is that, if someone is justified in believing a proposition, then she is also justified in believing any other proposition she knows to be entailed by the first. Call this claim the *Closure Principle for Justification.* In symbols:

1.1 (CJ) If J(X) and K(X ⇒ Y), then J(Y).[6]

CJ has as an instance:

1.2 If J(H) and K(H ⇒ ¬(E & ¬H)), then J(¬(E & ¬H)).

It is obvious that if H is true, (E & ¬H) is false. So, you know that H entails ¬(E & ¬H):

1.3 K(H ⇒ ¬(E & ¬H))

Therefore, if you are justified in believing H:

1.4 J(H)

the full antecedent of 1.2 is true, and you are justified in believing ¬(E & ¬H):

1.5 J(¬(E& ¬H)).

More simply, according to CJ, you aren't justified in believing a proposition unless you are justified in believing that your evidence for that proposition isn't misleading.

Now the trap closes. For the reasons we have just seen, it appears that you can't have justification for believing that your evidence for H isn't misleading. Hence, given CJ, you aren't justified in believing H. This argument is general, and it applies to any proposition for which you are supposed to have inductive justification. The outcome is skepticism with respect to inductive justification across the board.[7] Taking these points together, we face a trilemma. Either:

I. We accept skepticism.

II. We reject CJ.

or

III. We allow that one can have justification--somehow--for believing that one's evidence isn't misleading.[8]

In what follows, I will maintain that CJ can hold without making skepticism inevitable, because one can have justification for believing that one's evidence isn't misleading. In other words, the right horn of the trilemma to choose is (III).

2. Closure, theory choice and irrelevant conjunctions

One might think that our lack of justification for believing ¬(E & ¬H) does no harm if we abandon CJ. The problem of misleading evidence could then serve as the basis for an argument that CJ is unacceptable . However, such an argument can be extended, and ultimately it proves too much. What causes the trouble is the assumption that we have no justification for ¬(E &

¬H). Whether CJ holds is beside the point. Hence, considerations about misleading evidence provide no reason to reject CJ.

The challenge to CJ rests on the assumption that both the Entailment Principle and Hume's Principle are correct. There are, though, reasons to doubt that this dual assumption is true. First, consider two empirically adequate but conflicting scientific theories. Let's pretend that the Copernican hypothesis (CH) and the Ptolemaic hypothesis (PH) both entail (O) the observed motions of the planets. By Hume's Principle, we don't have a priori justification for ¬PH. If the Entailment Principle also holds, then O provides no empirical justification for ¬PH. Thus, we have no justification at all for ¬PH. This result is disturbing, at the very least. Quite plausibly, lacking justification for ¬PH would be just as bad as lacking justification for CH.[9] The same point may be put in other terms. If we aren't justified in believing ¬PH and similar claims, then a substantial form of skepticism prevails. This unhappy situation would be due to the combination of Hume's Principle and the Entailment Principle, independently of CJ. If we do have reason to reject PH, then either Hume's Principle, the Entailment Principle, or both must be wrong. Denying CJ wouldn't help with the problem at all.[10]

The use of the Entailment Principle and Hume's Principle against CJ encounters a further difficulty. There is a long-standing issue in confirmation theory known as the *problem of irrelevant conjunctions*.[11] Suppose that a "good" theory T entails evidence E, and that E justifies T. Let I be some far-fetched claim that has nothing at all do with E or T, and for which you have no additional evidence. In that case, (T & I) is an irrelevant conjunction. Here is an illustration. Suppose again that CH entails O. Let S = 'It will be sunny tomorrow on Cape Cod'. (CH & S) is an irrelevant conjunction. Since CH entails O, so does (CH & S). Insofar as CH entails O, O provides some reason to believe CH . But if that is right, we also have to say that information about the observed paths of the planets, O, provides reason to believe that the Copernican Hypothesis is true and that it will be sunny tomorrow on Cape Cod, (CH & S). It seems to many that something has gone wrong.

One standard approach to this problem is to allow that O does provide some relatively weak reason to believe (CH & S), while O provides a much stronger reason to believe CH alone.[12] Then, because O provides scant support for (CH & S), O doesn't justify accepting (CH & S). However, according to the Entailment Principle, O doesn't justify rejecting (CH & S), either. Since we lack justification for both accepting and rejecting (CH & S), presumably we ought to withhold with respect to (CH & S). But now consider the addition of more and more unrelated, extraneous claims to a "good" theory, producing longer and longer irrelevant conjunctions: $(T \& I_1)$, $(T \& I_1 \& I_2)$, $(T \& I_1 \& I_2 \& I_3)$, and so on. All of these conjunctions entail the evidence E. If the opponent of CJ is right, we must be justified in withholding with respect to any conjunction $(T \& I_1 \& I_2 \ldots \& I_n)$, no matter how long. In other words, we aren't justified in accepting $\neg(T \& I_1 \& I_2 \ldots \& I_n)$ regardless of how overwhelming likely it is to be true.[13] Giving up CJ doesn't avoid this unwelcome result.[14]

In short, using the possibility of misleading evidence to attack CJ relies on Hume's Principle and the Entailment Principle. However, the combination of those principles implies that we aren't justified in rejecting PH or in rejecting a very long irrelevant conjunction. These objectionable consequences follow whether CJ holds or not. The proper conclusion to draw is that either Hume's Principle or the Entailment Principle is false. But, in that event, there is no reason to deny that we have justification for $\neg(E \& \neg H)$, and the argument against CJ is blocked.[15]

3. E is evidence for ¬(E & ¬H).

The upshot of the previous section was that Hume's Principle and the Entailment Principle are jointly unacceptable. To that extent, we can have some kind of justification for believing that the evidence we have isn't misleading. One possibility is that we have a priori justification for believing ¬(E & ¬H); the other is that E itself provides empirical justification for rejecting (E & ¬H). The main drawback to endorsing the second alternative is that it is inconsistent with the Entailment Principle. But, in fact, there are cases in which E appears to justify rejecting a proposition that entails (E & ¬H), violating the Entailment Principle.

Consider the following examples:

> *Devil's Island*. Let A = No prisoner has ever escaped from Devil's Island before. Let B = Brittany (who is incarcerated on Devil's Island) will be the first prisoner to escape from there. It is plausible that A is a reason to reject to B, even though B entails A.
>
> *Seasons*. Let W = Winter has always been followed by spring. Let S = This winter will be the first not to be followed by spring. It is plausible that W is a reason to reject to S, even though W entails S.
>
> *Emeralds.* Let O = All observed emeralds are green. Let G = Even though all observed emeralds are green, there is at least one unobserved non-green emerald somewhere. It is plausible that O is a reason to reject to G, even though G entails O.

Let's look more closely at the Devil's Island Example, in particular. Presumably, (A) 'No prisoner has ever escaped from Devil's Island before' is evidence for (B*) 'Brittany won't escape from Devil's Island'. If A is *misleading* evidence in this instance, then (A & ¬B*) holds. (A & ¬B*) = 'No prisoner has escaped from Devil's Island before & Brittany will escape from Devils's

Island'. Put more simply, (A & ¬B*) is equivalent to 'Brittany will be the first prisoner to escape from Devil's Island'. To deny that A is misleading evidence for B* is to affirm ¬(A & ¬B*). ¬(A & ¬B*) = 'Brittany won't be the first prisoner to escape from Devil's Island' = B. Many people find it intuitive that A is reason to believe B. If that reaction is correct, then A itself is reason to believe that A isn't misleading evidence for B*.

It is worth noting that this intuitive impression is supported by certain formal considerations. One major approach to defeasible inference is default logic, due in its original form to Raymond Reiter.[16] Default logic may be viewed as ordinary sentential logic supplemented by an apparatus to handle *default rules*. Suppose again that A justifies accepting B*, even though A doesn't entail B*. Then, there is a default rule which licenses the transition from A to B*, so long as there are no defeaters for the justification that A* otherwise confers on B*. We may say that, in the absence of such defeaters, B* is justified simpliciter. And since B* is justified, so is ¬(A* & ¬B*), which may be derived from B*. Now, ¬(A* & ¬B*) would not count as justified without the membership of A in the stock of propositions one is initially justified in believing. Accordingly, it makes sense to view ¬(A* & ¬B*) as justified empirically, not a priori.[17]

There are further grounds for holding that our justification for ¬(E & ¬H) is empirical rather than a priori. John Hawthorne has advanced the view that we have a priori justification and a priori knowledge for the material conditionals which capture claims that one's evidence isn't misleading:

> Consider a character whom I will call 'The Explainer'. The Explainer has not had any experiences yet, but anticipates various experiential life histories,H1, H2, H3

> .... She also conceives of various theoriesT1, T2, T3 .... that describe possible structures of microphysical reality. The Explainer thinks about which theories about the world would be reasonable to believe given various possible experiential lives, being guided by considerations about which theory would provide the best explanation for each experiential life under consideration. Through such a priori research, the Explainer comes to believe propositions of the following form: Tn is the best explanation of experiential life history Hm...By using this method, the Explainer comes to believe a host of deeply contingent material conditionals. Insofar as one thinks that inference to the best explanation is a rational guide to belief in a theory, then it seems that the Explainer's beliefs are on the one hand, independent of perceptual knowledge, and on the other, eminently rational, and so pretty good prima facie candidates for knowledge.

By Hawthorne's lights,:

3.1 APK $(E \supset H)$

But, as Hawthorne himself seems to acknowledge, the justification we have for $(E \supset H)$ is defeasible.[18] That is, there is some proposition D such that (E & D) doesn't justify H, and even, we may assume, such that (E & D) justifies ¬H. If Hawthorne is correct, this evidential relation ought to be something we can appreciate a priori:

3.2 APK $((E \,\&\, D) \supset \neg H)$

But:

3.3 $(E \supset H) \,\&\, ((E \,\&\, D) \supset \neg H) \Rightarrow \neg(E \,\&\, D)$.

Then, on plausible assumptions:[19]

3.4 APK(¬(E & D)).

An example will help bring out why we ought to demur at this point.

Ulcers. Around the middle of the 20th century, there was a consensus among medical researchers that the cause of peptic ulcers was (A) hyperacidity due to stress and diet rather than infection. Some of the evidence (U) for the prevailing view was that the observed lesions looked like the result of self-digestion, that symptoms were alleviated by the administration of antacid agents, and that the highly acidity of the stomach made it an inhospitable environment for pathogens. Nevertheless, XXXX and XXXX discovered (U+) that ulcers could be cured by antibiotics, and that the bacterium *helicobactor pylori* in particular was a primary cause of the development of ulcers. In light of this new evidence, the medical community became justified in believing (¬A) ulcers weren't brought on by stress and diet.

Let's suppose that, given only U as evidence, the medical community would be justified in believing A. Following Hawthorne, we have:

3.5 APK (U ⊃ A).

Also, we may assume that, given both U and U+, medical experts would be justified in rejecting A. According to Hawthorne's view:

3.6 APK ((U & U+) ⊃ ¬A).

However, 3.5 and 3.6 entail:

3.7 APK¬(U & U+).

If 3.7 is true, the Explainer knows a priori that it isn't the case that: Stomach lesions have certain appearance, the acidity of the stomach makes it generally inhospitable to bacteria, etc., and that ulcers can cured by taking antibiotics. But, of course, the Explainer knows no such thing. ¬(U & U+) is false, so not known.

A plausible retreat would be to claim that the Explainer has a priori justification, if not a priori knowledge, with respect to the material conditionals of interest. But a version of the problem confronting Hawthorne remains. Material conditionals are monotonic, and permit strengthening of the antecedent. Consider a further example.

Trees. Suppose that the only kinds of trees in the forest are elms, oaks, and maples, that there are more or less equal numbers of each, and that we are justified in believing these things. The different kinds of trees can be distinguished from one another by the characteristics of their leaves. Elms and oaks have alternating leaves, while maples don't. Elms and maples have simple leaves, while oaks don't. And oaks and maples have lobed leaves, while elms don't. These facts are displayed in the following table:

LEAVES	Elm	Oak	Maple
Alternating?	Yes	Yes	No
Simple?	Yes	No	Yes
Lobed?	No	Yes	Yes

Let Trevor be a tree in the forest. Given the preponderance of trees that have alternating leaves, that Trevor is a tree in the forest is reason to think that Trevor has alternating leaves. Given the preponderance of trees that have simple leaves, that Trevor is a tree in the forest is reason to think that Trevor has simple leaves. And, given the preponderance of trees that have lobed leaves, that Trevor is a tree in the forest is reason to think that Trevor has lobed leaves. The view under examination is that, when E is evidence for H, we have a priori justification for (E ⊃ H). If that is true, then:

3.8 J(Trevor is a tree in the forest ⊃ Trevor has alternating leaves).

3.9 J((Trevor is a tree in the forest & Trevor is a maple) ⊃ Trevor has alternating leaves)).

3.10 J(Trevor is a maple ⊃ Trevor has alternating leaves).

But:

3.11 J(Trevor is a maple ⊃ Trevor doesn't have alternating leaves).

Assuming that our justified beliefs have to be consistent, 3.10 and 3.11 imply:

3.12 J(Trevor isn't a maple).

And, similarly, we arrive at:

3.13 J(Trevor isn't an oak)

and:

3.14 J(Trevor isn't an elm).

Putting these points together leads to an unattractive consequence. You are justified in believing that Trevor is either a maple, an oak, or an elm, but you are also justified in believing that Trevor isn't a maple, that Trevor isn't an oak, and that Trevor isn't an elm. Philosophers who favor a particular treatment of the Lottery Paradox and related phenomena may not be fazed by this outcome. Their view would be that your justified beliefs in this case aren't explicitly inconsistent unless justification obeys an agglomeration principle, and they reject such a principle. However, those who aren't inured to this approach will be troubled by the idea that you are justified in believing that Trevor is either an elm, an oak, or a maple, yet at the same time 3.12, 3.13, and 3.14 are true. This overall conclusion follows from assuming that there is a priori justification for believing that one's evidence is misleading, i.e., APJ(¬(E & ¬H)). Moreover, if justification obeys a agglomeration rule, then that the view that we have a priori justification for believing that our evidence isn't misleading leads to a contradiction.

One might respond that whatever goes wrong in this instance is really independent of whether we have a priori justification for (E ⊃ H). Everyone is stuck with the material conditionals like 3.8, and whatever trouble they generate down the line. Why so? Suppose we have as evidence that Trevor is a tree in the forest. Given our background information:

3.15 Trevor is a tree in the forest

is evidence for:

3.16 Trevor has alternating leaves.

Since we have evidence for 3.16:

3.17 J(Trevor has alternating leaves).

If justification obeys CJ, 3.17 implies:

3.8 J(Trevor is a tree in the forest ⊃ Trevor has alternating leaves).

Thus, it seems that we inevitably arrive at the unwholesome trio 3.12, 3.13, and 3.14. Whether there is a priori justification for believing that one's evidence isn't misleading doesn't matter. Hence, that view in particular isn't impugned by the argument above.

However, there is an oversight in this reply. A plausible view is that non-deductive inferences can be subject to collective defeat.[20] The idea, roughly, is this: Grant that an inference from E to H_1 is all right if taken in isolation. Still, it may be that further, otherwise legitimate inferences from E lead to conclusions H_2, H_3, ... H_n, which together are incompatible with H_1. The inference from E to H_1 suffers defeat under such circumstances, and possession of E as evidence doesn't justify (H_1). This possibility is important, because collective defeat occurs in the Trees Case. The inference from 3.14 to 3.15 is defeated, precisely because it leads to accepting the material conditional contained in 3.8 and its kin, and, then, to the unwanted implications which follow from these. Taking the possibility of collective defeat into account, the Trees Case does pose a special difficulty for the view that there is a priori justification for believing that your evidence is misleading. On that view, 3.8 holds, up front. But on the alternative picture which provides for collective defeat, 3.8 is false, and the difficulties which would ensue are avoided.

It may be useful to review the arguments presented in this section. The Devil's Island example and others like it suggest that E can be evidence that E itself isn't misleading. Certain formal considerations support such an assessment. But someone might deny that, in the cases of interest, E provides (empirical) justification for denying that E itself is misleading, maintaining that we have a priori justification instead. However, there is a drawback to adopting this account. If it were correct, you would also have justification for incompatible propositions about what kind of tree Trevor is. Insofar as this result is unacceptable, we don't generally have a priori justification for $\neg$(E & $\neg$H). In order to maintain that we do have justification for $\neg$(E & $\neg$H) as needed, we must allow that such justification can be empirical, and reject the Entailment Principle.

4. Theory choice, again.

There is another reason to disavow the Entailment Principle, which builds on some ideas that were broached in §§2 and 3. Suppose you are choosing between two competing hypotheses H1 and H2, both of which entail your evidence E. It may be that H1 and H2 are both supported by E to some extent, yet E supports H1 much more strongly than it supports H2. In such circumstances, I propose, E may justify rejecting H2 in favor of H. This is so despite the fact that

H2 entails E, which means that the Entailment Principle is violated. This view of theory choice bears immediately on the nature of our justification for ¬(E &¬H). (E & ¬H) and H are competitors. It may be allowed that both are supported to some extent by E. But if E supports H more strongly than E supports (E & ¬H), then E will justify ¬(E &¬H) all the same.

To get clearer about these issues, it will help to look at a toy example. Here is one, adapted from a case presented by Bas Van Fraassen[21]:

> *Mouse in the Wainscoting*. C = When you leave a bit of cheese by the little hole in the wainscoting at night, it's gone next the morning. W1 = There is a mouse in the wainscoting that, during the night, eats the cheese you left out by the hole. W2 = Your neighbor runs an exterminating company, and he sneaks into your house after you go to bed and removes the cheese by the hole. He does this in order to drum up some business for his firm.

My view of things is that, while C provides some reason to believe W2, C provides more reason to believe W1. Thus, C establishes W1. If W1 is true, W2 is false, so C also gives you reason to reject W2. In other words, C (empirically) justifies you in rejecting W2, despite the fact that W2 entails C. The upshot is that the Entailment Principle fails to hold in this case and, presumably, in others like it. Of course, this analysis is open to dispute. The opposing view is, again, that the Entailment Principle holds and that what I have taken to be empirical justification is really a priori justification. Since W2 entails C, the Entailment Principle implies that C can't justify ¬W2. So, if you are justified in believing ¬W2 at all, your justification must be a priori.[22]

This position is supported by a line of thought that has been quite influential. The following version is due to Michael Huemer:

> But now consider examples of likely candidates for reasons for preferring one hypothesis over another. Simplicity is often suggested in this connection -- that is, the fact that h_1 is the 'simpler' hypothesis in some sense may be a reason for preferring h_1 over h_2...But if h_1 is simpler than h_2, then it is a necessary truth that h_1 is simpler than h_2 (assuming that simplicity is an objective characteristic of propositions)... Now, if we take this route, we will get an interesting result. If the relevant necessary truth(s) can be known a priori, then it appears that we can also have a priori knowledge of (or at least justification for) synthetic, contingent truths. For if e is a reason for preferring h_1 over h_2, then it appears to be a reason for thinking that if either h_1 or h_2 is the case, h_1 is the case. Now the proposition, if either h_1 or h_2 is the case, h_1 is the case, is contingent, but it is apparently justified by an a priori truth. And whatever is justified by an a priori truth is justified a priori. Thus, there are contingent, a priori justified beliefs. Huemer (2001, p. 389).[23]

We need to examine this argument in detail. Various analyses are possible, but this one seems to capture the gist:

Huemer's Argument

4.1 If, given E, one is justified in accepting H1 and rejecting H2, then H1 is epistemically preferable to H2.

4.2 If H1 is epistemically preferable to H2, then, necessarily, H1 is epistemically preferable to H2.

4.3 If, necessarily, H1 is epistemically preferable to H2, then APJ(H1 is preferable to H2).

From 4.1, 4.2, and 4. 3:

4.4 If, given E, one is justified in accepting H1 and rejecting H2, then APJ(H1 is epistemically preferable to H2).

By assumption:

4.5 (H1 is epistemically preferable to H2) ⇒ ((H1 ∨H2) ⊃ H1).[24]

From 4.4 and 4.5:

4.6 If, given E, one is justified in accepting H1 and rejecting H2, then APJ((H1 ∨H2) ⊃ H1).

The import of 4.6 is that if one is justified in rejecting H2, one's justification must be a priori (at least in part), not empirical. 4.6 runs counter to the analysis of theory choice set out above, according to which one's evidence provides (fully) empirical justification for rejecting H2.

There are two reasons to doubt the cogency of the argument, however. First, there is a problem with the treatment of epistemic preferability in premise 4.5. Note that (H1 ∨H2) ⊃ H1) is equivalent to (¬H2 ∨ H1). Further, since H1 is incompatible with H2, H1 entails ¬H2, and (¬H2 ∨ H1) entails ¬H2. Therefore, 4.5 implies:

4.7 (H1 is epistemically preferable to H2) only if (¬H2).

4.7 seems mistaken, insofar as one might well have reason to believe H1 rather than H2, even though H2 happens to be true.

One might try to avoid this outcome by adopting a different account of epistemic preferability:

4.8 (H1 is epistemically preferable to H2) only if J((H1 ∨H2) ⊃ H1).

In virtue of the logical facts set out above, and assuming CJ:

4.9 (H1 is epistemically preferable to H2) ⇒ J(¬H2).

From 4.4 and 4.9:

4.10 If, given E, one is justified in accepting H1 and rejecting H2, then APJ(J(¬H2)).

But 4.10 falls short of:

4.11 If, given E, one is justified in accepting H1 and rejecting H2, then APJ(¬H2).

It is hard to say what bearing, if any, 4.10 by itself would have on the thesis that our justification for rejecting H2 is empirical.

There is a worse difficulty for Huemer's Argument, in any case. Premise 4.2 is pretty clearly false. It may be that:

4.12 If you have E as your evidence, then H1 is epistemically preferable to H2

and even:

4.13 □ (If you have E as your evidence, then H1 is epistemically preferable to H2).

But it would be a modal fallacy to infer from 4.13 to:

4.2 If you have E as your evidence, then, □ (H1 is epistemically preferable to H2).

Consider the gloss of 'epistemically preferable' set out in 4.8 and 4.9. If that is correct then:

4.14 □ (H1 is epistemically preferable to H2)

entails

4.15 □ (J(¬H2)).

But 4.15 is false. Recall the Mouse-in-the-Wainscoting Case. Suppose that, instead of having C as your evidence, your evidence were C+. C+ is the same as C, but also includes footprints leading to your neighbor's door, a dropped letter from a bank to an exterminating business at your neighbor's address demanding payment for an outstanding loan, and so forth. If C+ were your evidence, you might well be justified in believing W2. But according to 4.15, it is necessarily true that you are justified in rejecting W2. Hence, 4.15 must be false, and so must 4.2. More broadly, the argument 4.1-4.6 as a whole is unsound.[25]

Perhaps someone will think that invoking 4.2 was a mistake that can be bypassed. Even without 4.2, it might be argued that if two hypotheses entail the evidence, then one's empirical evidence alone doesn't justify accepting one and rejecting the other. The first premise of this reworked argument would be:

4.16 The total relevant evidence E justifies (H1 $\vee$ H2), but not H1.

4.16 implies that E alone is insufficient to secure the justification of H1. Therefore, if you are justified in believing H1, you must be justified in believing something in addition to E, call it "X".

4.17 J(H1) $\supset$ J(X).

Furthermore, since E by itself doesn't support H1, it can't justify X. So, if X is justified, there must be a priori justification for X:

4.18 J(X) $\supset$ APJ(X).

From 4.17 and 4.18:

4.19 J(H1) $\supset$ APJ(X).

The general conclusion of the argument is that when two theories entail the available evidence, and we are justified in accepting one and rejecting the other, the justification we have must be a priori (either wholly or in part). In that event, we aren't able to reject either hypothesis (solely) on the basis of the empirical evidence alone.

However, this account as it stands is ill-motivated. There is simply no reason to accept the primary assumption, 4.16.[26] Essentially, there are two different models of how, given E as our evidence, we could be justified in accepting H1 (and rejecting H2). These models may be displayed graphically, where the arrows indicate direct epistemic support:

4.20 E ➙ (H1 ∨ H2)
➘
H1
➚
X

4.21 E ➙ H1

4.16 declares that E alone doesn't provide sufficient reason to accept H1 (say, by way of inference to the best explanation). That is, 4.16 stipulates that the justification for H1 can't have the structure displayed in 4.21 and must have structure displayed in 4.20. But no grounds supporting this claim have been provided, and there is reason to doubt it. Suppose the epistemic norm applying to cases like Mouse in the Wainscoting example (and other less fanciful ones) is:

4.22 Where E is the pertinent evidence, H1 and H2 both entail E, and H1 provides a significantly better explanation of E than H2 does, then one is justified in accepting H1.

If 4.22 is the norm governing the choice between H1 and H2, then epistemic support will have the structure 4.21, not 4.20. That is, E justifies H1, and a priori justification for some further proposition is unnecessary. It seems quite plausible that 4.22, or something like it, is an operative epistemic norm. Absent any argument to the contrary, the status of 4.16 is very much in doubt.

Still, someone might insist that 4.16 is correct, appealing to a case like the following to make the point:

Diagnosis Case. A patient exhibiting symptom S goes to see the doctor. There are only two diseases, D1 and D2, which produce S, and both invariably do so. However, D2 is much rarer than D1. So, in nearly all cases, if someone has S, he has D1. The doctor is aware of this fact and concludes that her patient has D1, not D2.

Let's say that both the claim that the patient has D1 and the claim that he has D2 entail that the patient exhibits symptom S. It is natural to say in this instance that the patient's having S doesn't justify the doctor in believing that the patient suffers from D1 rather than D2. What does so is the doctor's belief that (F) instances of D2 (i.e., instances of D2 & S) are much less frequent than instances of D1 (i.e, instances of D1 & S). The doctor's justification for the diagnosis D1 goes beyond, or is other than, whatever justification is provided by S itself. To put the matter graphically, the structure of justification for the doctor is 4.22 rather than 4.23:

4.22 S → (D1 ∨ D2)
↘
D1
↗
F

4.23 S → D1

A defender of the Entailment Principle might go on to insist that the pattern of justification in the Diagnosis Case is the proper model for all situations in which two competing hypotheses entail all the relevant evidence.

Certainly, the structure of justification in *some* circumstances will conform to the pattern of 4.22 rather than 4.23. But there is a crucial difference between the Diagnosis Case and situations which bear on the status of the Entailment Principle. In a word, the Diagnosis Case isn't one in which the competing hypotheses entail all the relevant evidence. Part of that evidence is:

4.24 5% of the time, people with S have D2 rather than D1.

The hypothesis that, in this particular case, the patient has D1 certainly does not entail 4.24. Or, to put the same thing differently, the Diagnosis Case is not one which fits the schema:

4.16 The total relevant evidence E justifies (H1 ∨ H2), but not H1.

In the Diagnosis Case, the total relevant evidence *does* support D1. All in all, the Diagnosis Case provides no backing for 4.16 and, hence, has no bearing on the status of the Entailment Principle.[27]

To summarize: I have maintained that, even when two hypotheses entail all the relevant evidence, that evidence can still provide empirical justification for accepting one hypothesis and rejecting the other. If that is right, the Entailment Principle is false. Huemer's Argument (4.1-4.6) was meant to show that the justification for such a choice can't be (wholly) empirical. However, both that line of thought and its reworking (4.16-4.19) are flawed. The Diagnosis Case does nothing to strengthen the latter. Therefore, there is no basis for supposing that when we choose between empirically adequate hypotheses, we inevitably rely on some a priori justified belief. Our justification for accepting one hypothesis and rejecting the other may be fully empirical, in violation of the Entailment Principle. In particular, E may provide a reason to accept H and to reject (E & ¬H), so that we have empirical justification for believing that our evidence isn't misleading.

5. The Confirmation Principle and the dynamics of justified belief

The burden of §§3-4 was that the Entailment Principle is false and that E can provide justification for ¬(E & ¬H). Since the Confirmation Principle implies the Entailment Principle, the results of §3 and §4 count against the Confirmation Principle, too. But the Confirmation Principle fails for a further reason, to be explored below. As a related matter, there is a prominent argument which proceeds from considerations about the acquisition of justified belief to the conclusion that E can't justify ¬(E & ¬H). The difficulty that arises for the Confirmation Principle vitiates this argument as well.

Many regard the Confirmation Principle as a truism. The content of the principle is that evidence justifies a hypothesis only if that evidence makes the hypothesis more likely to be true. A consequence of the principle is that, if one acquires evidence E, then E justifies H only if the acquisition of E increases the strength of one's justification for believing H. But this claim can't be correct in general. To put the matter most simply, suppose you originally have evidence E1 which makes H extremely likely, and thus justifies H. You then obtain further evidence E2 which defeats (expunges) the justification for H provided by E1. At the same time, E2 independently makes H very likely, although a bit less likely than E1 did. Since E2 makes H very likely, E2 may justify you in believing H, even though the probability of H has decreased. In such circumstances, the Confirmation Principle fails.

Here is an illustration of the point:

> *Thermometers*. Rex is inside an air-conditioned room. He looks out the window at his deluxe model thermometer, which reads 80 degrees. Given Rex's background information, the reading of the deluxe thermometer makes it extremely likely that the outside temperature is 80 degrees. The deluxe model thermometer is slightly more reliable than the standard model, so in general one ought to place slightly less confidence in the readings of the standard model than in the readings of the deluxe model. Still, the standard model is an excellent instrument, used for many demanding applications. (We may suppose that if Rex had used a standard model thermometer instead of the deluxe one, he would still

> have been justified in believing that the temperature is 80 degrees). Now, it happens that Omar, the thermometer repair man, comes over and checks Rex's deluxe thermometer. Omar discovers that, extraordinarily, Rex's deluxe thermometer isn't working properly. Fortunately, however, Omar is carrying a standard model thermometer which reads 80 degrees. Omar comes in and tells Rex that his deluxe thermometer is broken, but also that Omar's own standard thermometer reads 80 degrees. Given the new evidence provided by Omar, Rex ought to be slightly less confident than he was before that the outside temperature is 80 degrees. Nevertheless, what Omar tells Rex justifies him in believing that the outside temperature is 80 degrees.

If the Confirmation Principle were correct, the upshot of Omar's conversation with Rex would be that Rex isn't be justified in believing that the temperature is 80 degrees. Since that isn't so, the Confirmation Principle is incorrect.

The defender of the Confirmation Principle might resist in various ways. One thing he might say is that, in the case as described, Rex's rational confidence that it's 80 degrees out drops when Omar tells him that the deluxe thermometer is broken. Rex's confidence then goes back up when Omar adds further information about the reading of the standard thermometer. The further information justifies Rex in believing that the temperature is 80 degrees insofar as it raises Rex's rational confidence from its previously lowered level. Another idea is that the description of the case oversimplifies the way justification works, and thus how the Confirmation Principle governs justification. Confirmation and justification obtain with respect to a background or default degree of rational confidence (a default probability function) which may differ from the degree of rational confidence one actually has at a particular time.[28] Omar's information does lower Rex's rational confidence with respect to its previous actual level. However, the same information elevates Rex's rational confidence that the temperature is 80 degrees vis-a-vis its background or default value, and that is why Rex's belief can be justified. Thus, the Confirmation Principle holds, but with respect to background or default levels of rational belief, specifically.

I doubt that either of these responses preserves the Confirmation Principle. But instead of pursuing this issue right away, let's consider an argument that might be offered in support of the principle. It will become apparent in this context why the Confirmation Principle can't be saved in either of the ways just described. The argument in favor of the Confirmation Principle is suggested by a line of thought Roger White has directed against dogmatist replies to skepticism.[29] Both focus on the "dynamics" of justified belief, i.e., how we come to believe H on the basis of evidence E. So I will call the argument meant to bolster the Confirmation Principle the "Dynamical Argument".

The Dynamical Argument is directed against the view I have been defending, namely, that when you have E as your evidence for H, E can justify you in believing that E itself isn't misleading. The Dynamical Argument proceeds within a Bayesian setting. We assume that rational credences at a time are probabilities, and these credences are updated by conditionalization. Beyond this basic Bayesian framework, we also suppose that if one's rational credence in a proposition is sufficiently high, then one is justified in believing that proposition. The Dynamical Argument goes as follows:

Dynamical Argument

5.1 Suppose you acquire E at t2, and thereby come to be justified in believing H for the first time.[30] (Assumption).

5.2 Your belief that H is arrived at by conditionalization, so that Pr(H) at t2 = Pr(H/E) at t1. (Assumption).

5.3 Since Pr(H) at t2 is equal to or greater than the threshold for justification, so too is

Pr(H/E) at t1. (From 5.1 and 5.2)

5.4 Pr(¬(E & ¬H)) at t1 is greater than or equal to Pr(H/E) at t1. (Probability fact).

5.5 Pr(¬(E & ¬H)) at t1 is greater than the threshold for justification (so at t1 you are justified in believing ¬(E & ¬H)). (From 5.3 and 5.4).

5.6 At t1, you have yet to acquire E, so at t1 E doesn't justify ¬(E & ¬H). (Assumption).

5.7 Acquiring E doesn’t change the nature of the justification you have for ¬(E & ¬H), if any. (Assumption).[31]

5.8 Hence, even when you acquire E as your evidence, your justification for ¬(E & ¬H) isn’t due to E. (From 5.5, 5.6, and 5.7).

The Dynamical Argument might be resisted to the extent that it relies on aspects of Bayesianism that are controversial.[32] But suppose we grant the correctness of the argument through 5.6. Even so, the argument doesn't go through to its conclusion. The trouble lies with premise 5.7. When you acquire E the nature of your justification for believing ¬(E & ¬H), if you had any, will change. The idea is that E removes the reason you had beforehand to believe ¬(E & ¬H), while also giving you different justification for ¬(E & ¬H) in its place. So, when you acquire E, E (empirically) justifies your belief that ¬(E & ¬H) after all.

Here is the point in detail. Consider your epistemic situation at t1. Pr(E & ¬H) at t1 is high. (E & ¬H) may be rewritten as (¬E ∨ H). Thus, at t1, Pr(¬E ∨ H) is high. It is a fact of the probability calculus that Pr (¬E ∨ H) = (Pr(¬E) + Pr(H) - Pr(¬E & H)). Therefore, Pr(¬E ∨ H) is high at t1 only if the sum (Pr(¬E) + Pr(H)) is high at t1. By hypothesis, at t1, you don't have a justified belief in H, so at t1 Pr(H) is low. Hence, if (Pr(¬E) + Pr(H)) is high at t1, that is because

Pr(¬E) is high at t1.[33] So, at t1, you are justified in believing (¬E ∨ H) to the extent that you are justified in believing ¬E. We may allow for the sake of the argument that, at t1, your justification for ¬E, and thus for (¬E ∨ H) is non-empirical. But, then, you have non-empirical justification for (¬E ∨ H) only insofar as you have justification for ¬E.

Now suppose that you acquire E as evidence, and update. Whatever justification you had for ¬E is lost, and with it you lose your original justification for (¬E ∨ H). But also, when you update on E, your probability for H, which was low, goes up past the threshold for justification. And, since Pr(H) is high at t2, Pr (¬E ∨ H) must be high at t2 as well. That is to say, at t2 you continue to be justified in believing (¬E ∨ H), as you were at t1. However, *the source of your justification changes*, in a way that may be masked by the fact that Pr(¬E ∨ H) is high at both t1and t2.[34] Upon acquiring E, the job of justifying (¬E ∨ H) is "handed off" from ¬E to E. The result is that, when you do have E as your evidence, E justifies you empirically in believing ¬(E & ¬H)--whatever may have been the case beforehand.[35]

The preceding discussion of the Dynamical Argument has been quite abstract. Therefore, working through an example may be helpful. Suppose that, before you embark on your investigation of emeralds, you have no empirical evidence as to their color. You examine many emeralds and acquire evidence O, that many emerald have been observed, all of which are green. O supports G, that all emeralds are green.[36] Let's grant, for the sake of the argument, that before you acquire O, you have justification for (¬O ∨ G). It seems out of the question to suggest that before acquiring O you are justified in believing G. In the absence of empirical information like O, your rational credence with respect to G ought to be pretty low, even though Pr(G/O) is high. So, if you are initially justified in believing (¬O ∨ G), that is because you have justification for

(¬O). You then examine lots of emeralds and get O as evidence. O defeats ¬O, which was your initial justification for (¬O ∨ G). But, at the same time, O justifies you in believing G. The upshot is that the Dynamical Argument fails to establish that E can't justify ¬(E & ¬H). If anything, the considerations brought to the fore by the argument support that claim.

Two further observations are called for. First, a Bayesian of a certain stripe might maintain that you are justified in believing ¬(E & ¬H) at t1. Insofar as you have no empirical justification for ¬(E & ¬H) at t1, by way of E or anything else, your justification for ¬(E & ¬H) at t1 must be a priori. Therefore, in order to have empirical justification for ¬(E & ¬H) at some time, we must have a priori justification for ¬(E & ¬H) at some other time. My response to the Dynamical Argument doesn't foreclose such a possibility. I haven't tried to show that a priori justification for ¬(E & ¬H) is impossible. But, to turn things around, it may well be that the details of Bayesianism plus other commitments make it attractive or necessary to claim that there can be a priori justification for ¬(E & ¬H) at some time or other. However, a proponent of the Dynamical Argument might have wanted a good deal more from it than that. First, she may have meant the argument to proceed from generally agreed upon considerations concerning justification, rather than those peculiar to Bayesian epistemology.[37] Her hope would be that, relying solely on uncontroversial assumptions, she could establish that E can't provide empirical justification for ¬(E & ¬H) under any circumstances; hence, on pain of skepticism, there must be must be a priori justification for believing ¬(E & ¬H). But, as I have maintained, the Dynamical Argument doesn't achieve that much.

The second observation to make is this. I have tried to show that the Confirmation Principle in full should be rejected. That is, justification and confirmation (i.e., probabilification) sometimes come apart. But one might think that, surely, confirmation is connected to justification somehow. That may be, although exactly how they are related seems to me to be an open question. Here is one proposal that I offer with some diffidence. Suppose, at least to a first approximation, that justification amounts to high epistemic probability on one's total evidence. Then

5.9 R confirms S

implies

5.10 If you are otherwise justified in believing S and you add R to your stock of evidence, you will remain justified in believing S.

and

5.11 R disconfirms S

implies

5.12 If you are otherwise justified in believing S, and you add R to your stock of evidence, you may not remain justified in believing S.

Whether 5.11 and 5.12 are exactly right is of secondary importance. What matters most is that abandoning the Confirmation Principle doesn't mean severing any link whatsoever between justification and confirmation.

Let's recall the mains points made in this section. If the Confirmation Principle holds, E can't justify ¬(E & ¬H). However, the Confirmation Principle fails when one's previous justification for a proposition is defeated by new evidence and that new evidence supports the target proposition, albeit at a slightly lower level. Precisely that may happen when the proposition is ¬(E & ¬H) and one acquires E as one's evidence.[38] The Dynamical Argument is meant to demonstrate that E can't be evidence for ¬(E & ¬H). However, that argument is defective for roughly the same reason that the Confirmation Principle is. All in all, these results promote the view that we can have empirical justification for denying that our evidence is misleading.

6. More on skepticism and misleading evidence

As we saw earlier, the possibility that one's evidence is misleading gives rise to a certain line of skeptical argument: Suppose that E is your evidence for H. By CJ, you're not justified in believing H unless you're justified in believing ¬(E & ¬H). The skeptic denies that you have either empirical or a priori justification for ¬(E & ¬H). Hence, you're not justified in believing H. A common view is that many or all familiar skeptical challenges are really just variations on this line of argument.[39]

In the first place, Cartesian skepticism depends on the possibility that your experience might be thoroughly unveridical. Take an example. Suppose (A) it appears to you that there's a sand dune in front of you, but (¬D) there really is no such thing, because you are the victim of massive sensory deception. If your experience is unveridical in that way, then (A & ¬D). One might construe A as your sensory evidence for D. In that case, the possibility that your experience is unveridical amounts to the possibility that your sensory evidence is misleading. If in general you aren't justified in believing that your evidence isn't misleading, then you aren't justified in believing that your evidence is veridical. Cartesian kepticism prevails. Seen this way, Cartesian skepticism is just a specific version of the problem of misleading evidence.

Consider next the problem of induction. Suppose that you believe (G), all emeralds are green, on the basis of the evidence (O), all observed emeralds are green. One version of inductive skepticism is to attack directly the claim that O provides any reason to accept G. But another way for the skeptic to proceed is to allow, for the sake of argument, that O *is* evidence for G. Even so, the skeptic will say, your evidence could be misleading. It could be the case that all observed emeralds are green, yet not all emeralds are green (perhaps all unobserved emeralds are blue). If so, (O & ¬G). The skeptic will then argue via the Entailment Principle that O provides no evidence against (O & ¬G). To that extent, ¬J(¬ (O & ¬G)), and, by way of CJ, ¬J(G). You don't really have inductive justification for believing G, when all is said and done.[40]

I would like to make two points, one about this second version of inductive skepticism, the other about Cartesian skepticism. With regard to induction, it is clear that the skeptical relies on the Entailment Principle. But that principle is false, or so I have argued. O can be reason to believe ¬(O & ¬G), so that J¬((O & ¬G)). Because the argument for the second form of inductive skepticism requires denying J¬((O & ¬G)), that argument is unsound. Admittedly, this reply presupposes that O is evidence for G. It therefore carries no weight with respect to the first version of inductive skepticism described above. There is some gain, nevertheless. If the first version of inductive skepticism can be dealt with, then the second version poses no further difficulty, contrary to what one might have supposed.

Turning now to Cartesian skepticism, one prominent anti-skeptical position is dogmatism. The dogmatist holds that your experience EXP (say the experience as of seeing a hand) gives you justification for believing (HAND) that there is a hand before you. Either directly, or by way of your justified belief that HAND, EXP is also evidence for the claim (¬SK) that you aren't the victim of massive sensory deception.[41] An influential criticism of dogmatism, worked out in detail by Roger White, challenges it on this score. If SK is true, then it appears to you that there is a hand before you even though there really isn't. This would be a situation in which you have misleading evidence, i.e., (EXP & ¬HAND). If the Entailment Principle holds, then EXP can't be evidence for ¬(EXP & ¬HAND). That is, EXP can't be evidence for ¬SK, contrary to what the dogmatist maintains.[42] However, this objection to dogmatism is by no means conclusive. It depends squarely on the Entailment Principle, but, as I have tried to show, the Entailment Principle is unacceptable. To that extent, dogmatism emerges unscathed.[43]

7. Conclusion

The problem of misleading evidence is connected to a number of important issues in epistemology. Among these are the status of the closure principle for justification, the workings of theory choice, and the fortunes of various kinds of skepticism. How we ought to regard the problem of misleading evidence depends on the standing of the Entailment Principle and the Confirmation Principle. If what I have said here is right, then neither principle turns out to be well-motivated. Without them in place, it seems that evidence for a hypothesis can also justify the claim that the evidence itself isn't misleading.[44]

Endnotes

1. Since your evidence is inductive, E doesn't entail ¬H, and (E & ¬H) is at least logically possible.

2. I have discussed various facets of the problem elsewhere. See Vogel (200?), (200?).

3. These are epistemic probabilities.

4. Proof: Assume X entails Y. (i) If so, Pr(Y/X) = 1. (ii) From (i), Pr(Y/X) Pr(X)/Pr(Y) = Pr(X)/Pr(Y). (iii) Assuming that Pr(Y) is such that 0 < Pr(Y) < 1, then Pr(X)/Pr(Y) > Pr(X). (iv) From (ii) and (iii), Pr(Y/X) Pr(X)/Pr(Y) > Pr(X). (v) According to Bayes Theorem, (X/Y) = Pr(Y/X) Pr(X)/Pr(Y). (vii) So, from (iv) and (v), Pr(X/Y) > Pr(X). (viii) From (vii), 1- Pr(X/Y) < 1- Pr(X). (ix) And from (viii), Pr(¬X/Y) < Pr(Y). Since Pr(¬X/Y) < Pr(Y), it's not the case that Pr(¬X/Y) > Pr(Y).

5. Examples of the contingent a priori proposed by Saul Kripke provide no reason to think that there is a priori justification for believing ¬(E & ¬H). Hume's Principle should be understood as restricted to "deeply contingent" propositions. For discussion of this issue, see John Hawthorne, "Deeply Contingent A Priori Knowledge" PPR, 63:2, 2002, p. 247-249 and Brian Weatherson, "Skepticism, Rationalism, and Empiricism" Oxford Studies in Epistemology, ed. T. Gendler and J. Hawthorne (Oxford: OUP, 2005), p. 311.

6. I use the following notation: 'J(X)' means that one is justified in believing X; 'X ⇒ Y' means that X entails Y; 'K(X)' means that one knows X; 'APJ(X)' means that one is justified a priori in believing X. Throughout, I will assume that the pertinent entailments are known by the subjects.

7. A number of philosophers think that the problem of misleading evidence is the common essence of all skeptical arguments. See Weatherson (2005), with some caveats, and Michael Huemer, "The Problem of Defeasible Justification" (2001). From this standpoint, there is no significant difference between skepticism about induction and skepticism about the external world, a view I reject "Skeptical Arguments"(2004). For further discussion, see §6, below.

8. Fred Dretske deploys something like this trilemma in his recent attack on the closure principle for knowledge 'The Case Against Closure" In M. Steup & Earnest Sosa (eds.), Contemporary Debates in Epistemology. Malden, Ma: Blackwell, 2005. Dretske's view is that we don't have justification for believing propositions like ¬(E & ¬H) and that embracing skepticism is out of the question. He writes, "If, in order to see (hence, know) that there are cookies in the jar, wine in the bottle, and a zebra in the pen, I have to know that I am not being fooled by a clever deception, that the 'appearances' (the facts on which my judgments are based) *are not misleading,* then skepticism is true." (2005, p. 16-17, emphasis added). The only way to preserve closure for knowledge, in the absence of closure for justification, would be to say that one can know a proposition like ¬(E & ¬H) without having any reason to believe it (or, as

Dretske puts it, “any way to know” it; 2005, p. 20), and he rejects such a view out of hand. I should add, though, that Dretske’s views about closure for justification and closure for knowledge are complex and multi-faceted. What I say here does not address the full range of his recent thought on the topic.

9. Refuting a hypothesis is, as such, no less important nor different in kind from establishing one. Accordingly, an experiment may be highly significant because it establishes that a particular theory is false, not that some other theory is true. The Michelson-Morley experiment was important because it overturned a version of the ether theory (although it wasn’t viewed that way at the time it was conducted). Or, to take another example, the Meselson-Stahl experiment refuted the accounts of DNA replication due to Delbrueck and to Stent, as much as it supported the Watson-Crick view. [For discussion, see Weber, Marcel: The Crux of Crucial Experiments. Duhem's Problems and Inference to the Best Explanation. The British Journal for the Philosophy of Science 60: 19-49 (2009).]

10. Dretske says that closure for knowledge may fail if X entails Y, but the way one knows X differs from the way one knows Y. Dretske calls such propositions Y “heavyweight propositions” (2005, 16). Now, with regard to the example in the text, your reason for accepting CH and your reason for rejecting PH are one and the same (let’s say it is that CH is simpler than PH in some appropriate way). You don’t know that CH is true in some way different from the way you know that PH is false. Hence, a Dretskean may not reject the argument in the text by claiming that PH is a “heavyweight” proposition while CH is not, so that we know (are justified in believing) CH but don’t know (aren’t justified in believing) ¬PH. But even if the Dretskean had some way of drawing such a distinction, he shouldn’t do it. If we know CH, then denying that we know ¬PH isn’t shrewd philosophy, it’s bad science.

11. See C. Glymour, Theory and Evidence, Princeton U. Press, 1980.

12. This treatment of irrelevant conjunctions may be supported by various probabilistic considerations. For references and (critical) discussion, see Patrick Maher, “Bayesians and Irrelevant Conjunction” Philosophy of Science 71, 2004 (p. 515-520).

13. It might seem that such a situation does arise in the case of lottery propositions. But $\neg(T \,\&\, I_1 \,\&\, I_2 \ldots \&\, I_n)$ isn’t a lottery proposition. We don’t have any reason to believe that one of the I_i is true.

14. You are justified in believing T, but T doesn’t entail $\neg(T \,\&\, I_1 \,\&\, I_2 \ldots \&\, I_n)$, so CJ doesn’t apply in this direction. However, I do think that you are justified in believing $\neg(I_1 \,\&\, I_2 \ldots \&\, I_n)$. If CJ holds, you are also justified in believing $\neg(T \,\&\, I_1 \,\&\, I_2 \ldots \&\, I_n)$. So, although the problem in the text arises even if one denies CJ, preserving CJ would ameliorate it.

15. In other words, the correct response to the trilemma of the previous section is to endorse alternative III rather than alternative II. However, it must be said that the argument in the text is a good deal less than a full defense of CJ. Many philosophers hold on formal and intuitive

grounds that evidence fails to obey a certain kind of closure condition. Let 'E(X)' mean that E justifies X. The closure principle (CE) is then: [E(X) & K(X ⇒ Y)] ⇒ E(Y). Variants of this principle have been discussed under the rubric of the Special Consequence Condition (Hempel) and the Transmission of Warrant (Wright). The threat is that if E is your evidence for X, justifying you in believing X, E might not be evidence for Y, leaving you with no justification for Y. In that event, failure of closure for justification would follow from failure of closure for evidence. The problem of misleading evidence may be seen as a special case of this phenomenon. Suppose that we don't have a priori justification for ¬(E & ¬H). Then CJ will fail if E is evidence for H but not for ¬(E & ¬H). That is, we would have a violation of CJ due to a failure of CE. The defense of CJ in the text doesn't apply straightforwardly to broader worries about CJ brought on by the possibility of violations of CE. I believe that CJ holds in general, but I can't pursue that issue here.

16. For a survey of non-monotonic reasoning that covers default logic, see G. Brewka, J. Dix, and K. Konolige, *Non-monotonic Reasoning–An Overview* (CSLI Publications: Stanford, 1997) see also John Horty *Defaults as Reasons* (in preparation). The story told in the text assumes that the default rules are "normal", in the technical sense. I am indebted to Horty for advice about default logic, although I am responsible for any defects in the argument in the text.

17. Insofar as these observations provide some basis for the claim that we have empirical justification for ¬(E & ¬H), they provide additional substantiation for the claim that we have *some* justification for ¬(E & ¬H).

18. Hawthorne (2002), p. 252.

19. These assumptions include what Hawthorne and others have called "multi-premise closure" and the supposition that the Explainer has a certain level of logical perspicacity.

20. The most prominent advocate of this approach is John Pollock. See, *inter alia*, his "Justification and Defeat" Artificial Intelligence Volume 67, Issue 2, June 1994, Pages 377-407. For objections, see Igor Douven and Timothy Williamson (2006). Generalizing the lottery paradox. British Journal for the Philosophy of Science 57 (4):755-779 and David Poole, "What the lottery paradox tells us about default reasoning" in *Proceedings of the first international conference on principles of knowledge representation and reasoning* (San Francisco: Morgan Freeman Publishers, 1989). For a response to Douven and Williamson, see Jake Chandler (2010). The Lottery Paradox Generalized? British Journal for the Philosophy of Science 61 (3):667-679.

21. Bas van Fraassen. The Scientific Image. Oxford University Press, 1980, p. 20-21.

22. The Mouse-in-the-Wainscoting example is nice and vivid, but it isn't fully apt as an illustration. W1 doesn't entail the evidence C. Also, it is natural to assume that C wouldn't be one's total relevant evidence. In ordinary situations, you would have background information which bears on how likely it is that your neighbor would be in such a situation and would act as

described.

23. Huemer explores the view that a priori justification plays a crucial role in theory choice, but he doesn't endorse it without qualification.

24. Take this relation to be strong equivalence, to permit substitution in the scope of the operator APJ(...).

25. This discussion sheds some light why the inference from 4.9 to 4.10 (above) would be unsound. The following argument, though not water-tight, is at least strongly suggestive. J((W1 ∨W2) ⊃ W1)) is true only if one's evidence isn't C+. And, so, plausibly:

J((W1 ∨W2) ⊃ W1)) ⇒ ¬J(C+).

Now suppose 4.9 is right:

APJ(J((W1 ∨ W2) ⊃ W1)).

Assuming closure for a priori justification:

APJ(¬J(¬C+)).

This seems wrong. Maybe there can be a priori grounds for favoring a better explanation over a worse one, *given* a certain body of evidence. But can you have a priori justification for believing that your evidence is one way rather than another?

26. Certainly, invoking the Entailment Principle at this juncture accomplishes nothing, if the argument 4.16-4.19 is supposed to provide independent support for that principle.

27. There is another way to argue that justified choice between empirically adequate hypotheses must have the structure displayed in 4.20 rather than in 4.21. Some philosophers believe that inference to the best explanation is illegitimate unless there is some super-added reason to believe that such inference is likely to yield true results. This necessary additional element would be the value for X in 4.20. I think that such a "truth-demand" carries no weight. For discussion, see Vogel XXXX.

28. This view was suggested to me by Hartry Field in conversation; see also Lange ???

29. Roger White: "Problems for Dogmatism", Philosophical Studies" (2006) 131: 525-55.

30. This restriction is meant to guarantee that when you do have justification for H, that is because E in particular justifies H. The question is then whether E also justifies ¬(E & ¬H).

31. This step could be supported by appeal to the Confirmation Principle or treated as independently credible. But, as I will maintain, it ought not to be accepted one way or the other.

32. White declares that his Bayesian commitments are "very modest" (2006, p. 535) and seeks to bolster their employment by way of various examples.

33. I am setting aside the special case where Pr(¬E) and Pr(H) are both middling at t1, such that their sum exceeds the threshold for justification at t1, even though neither Pr(¬E) nor Pr(H) does individually.

34. The broader point is that the Bayesian apparatus has trouble dealing with relations of epistemic priority, because facts about epistemic priority don't reduce to facts about conditional and unconditional epistemic probabilities. This difficulty bedevils attempts to analyze putative failures of "warrant transmission" in probabilistic terms. For a recent attempt, and a review of previous ones, see Luca Moretti (forthcoming) "Wright, Okasha and Chandler on transmission failure". Synthese.

35. This observation meets the objection that Rex's degree of belief goes down and then up. It is a theorem of the probability calculus that Pr(¬(E & ¬H)/E) has to be lower than Pr(E & ¬H). So, on a Bayesian model of justification, the level of justification for ¬(E & ¬H) provided by E can't be as high as the level of justification one had before acquiring E. But E may justify ¬(E & ¬H) nevertheless. The same point applies to the objection that, in the earlier example, the Confirmation Principle might be preserved because Rex's rational credence in ¬(E & ¬H) could deviate from that assigned by a background or default probability function, Suppose Rex's rational credences conform exactly to those assigned by that function. When Rex acquires E as evidence, his rational credence with respect to ¬(E & ¬H) must fall. Even so, E may justify Rex in believing ¬(E & ¬H).

36. Of course, this is a caricature, and some may doubt whether enumerative induction as such can provide justification at all. However, I think that the particular mechanism of inductive justification at work is really irrelevant; the situation would be structurally the same in any case. For that reason, the toy example in the text should raise no scruples.

37. See Note 32. The idea is that the Dynamical Argument would, fundamentally, rely only on the Confirmation Principle or some other putatively anodyne claims about justification. My impression is that someone who adopts default logic as a formal model of justification need not agree with a Bayesian who holds that a priori justification for ¬(E & ¬H) must be in place independently, if the acquisition of E is to provide justification for H.

38. This observation meets the objection that degree of belief goes down and then up. It is a theorem of the probability calculus that Pr(¬(E & ¬H)/E) has to be lower than Pr(E & ¬H). So, on a Bayesian model of justification, the level of justification for ¬(E & ¬H) provided by E can't be as high as the level of justification one had before acquiring E. But E may justify ¬(E & ¬H) nevertheless. The same point applies to the objection that, in the earlier example, the

Confirmation Principle might be preserved because Rex's rational credence in ¬(E & ¬H) could deviate from that assigned by a background or default probability function, Suppose Rex's rational credences conform exactly to those assigned by that function. When Rex acquires E as evidence, his rational credence with respect to ¬(E & ¬H) must fall. Even so, E may justify Rex in believing ¬(E & ¬H).

39. See Note 7.

40. The first skeptical problem may be thought of as a version of the "old riddle of induction", and the second may be thought of as a version of the "the new riddle". My approach to these issues is indebted to Gemes (XXXX).

41. My presentation of dogmatism deviates from Pryor's own. Pryor (2000) generally denies that one's experience is evidence for one's perceptual beliefs.

42. This criticism of dogmatism is a variant of the one due to Roger White, and others. White's criticism would also apply to the attempt to foil Cartesian skepticism by an appeal to inference to the best explanation. [See IRS, p. xxxx]

43. To my mind, dogmatism is untenable for other reasons. For an early statement of the objection that dogmatism is too strong see DSP, and for other objections see IRS.

44. I would like to thank Fred Dretske, Hartry Field, Branden Fitelson, John Horty, James Joyce, Christopher Meacham, Joshua Schecter, Robert Stalnaker, and Roger White for valuable discussions about the topics discussed here. David Christensen and Daniel Greco have been especially generous with their help, without which this paper might not have been written.

Appendix: Weatherson's Argument Regarding Misleading Evidence

Brian Weatherson (2005) has provided an important treatment of what is, in my terms, the problem of misleading evidence. He sets out an interesting argument which would, given certain commitments, show that one has a priori justification for believing that one's (total) evidence isn't misleading However, there are reasons to doubt that the argument succeeds.

Here is a modified and somewhat simplified version of Weatherson's line of thought.[1] Assume as background that E justifies believing H. Let J*(X) stand for 'I am justified in believing X', let APJ*(X) stand for 'I am justified a priori in believing X', and let 'ETE' stand for 'E is my total evidence'.

W1. $\Box$ [ETE $\vee$ ¬ETE].[2]

By assumption:

W2. ETE $\Rightarrow$ J*(H).

From W2 and CJ:

W3. ETE $\Rightarrow$ J*(¬ETE $\vee$ H).

Suppose that a form of internalism is true, so that the mental state I am in determines what my total evidence is. Given internalism, whether ETE is true or not is a fact about my current mental state. In addition, suppose that mental states are "luminous." Understand this second assumption to entail that, in general, if M is some fact about my current mental state, then J*(M). Then:

W4. ¬ETE $\Rightarrow$ J*(¬ETE).

From W4 and CJ:

W5. ¬ETE $\Rightarrow$ J*(¬ETE $\vee$ H).

From W1, W3, and W5:

W6. $\Box$ [J*(¬ETE $\vee$ H)].

Now, suppose that, necessarily, I am justified in believing some proposition. I will be justified in believing that proposition regardless of what empirical evidence I happen to have. To that extent, justification I have appears to be a priori. If so, (W6) implies:

W7. $APJ^* (\neg ETE \vee H)$.

Informally, (W7) says that I have a priori justification for believing that if E is *my* total evidence, then E isn't misleading.

However, something is amiss. Consider:

W8. $\Box\, [J^*(\neg(\text{E is } \textit{your} \text{ total evidence}) \vee H)]$.

I might have reason to believe that, while E is your total evidence, and justifies *you* in believing H, you are benighted, and H is false. Since such a situation is possible, (W8) is false. Hence, there is no basis for allowing that:

W9. $APJ^*[\neg(\text{E is } \textit{your} \text{ total evidence}) \vee H)]$.

This result indicates that there is some defect in the argument (W1)-W(7). How can it be that I have a priori justification for believing that E isn't misleading if E happens to be *my* total evidence (i.e., W7), but not if E happens to be *your* total evidence (W9)? How can the fact that the evidence is mine and not yours make any difference as to whether $\neg(E \,\&\, \neg H)$ is true?

But if something is wrong, it isn't easy to pinpoint exactly what it is. One plausible diagnosis is that the construal of luminosity which yields (W4) is mistaken. Suppose M is some fact about my current mental state. Perhaps I have special epistemic access to my mental states, so that *if I introspect*, I am justified in believing M. However, it seems wrong to claim that I am justified in believing M even if I don't introspect. Thus:

W10. $M \Rightarrow$ (If I introspect, then $J^*(M)$)

may be true, but

W11. $M \Rightarrow$ (If I don't introspect, then J*(M)).

is false.[3]

Let's see what happens to the argument above if the weaker, more appropriate version of luminosity is employed. As before:

W12. $\Box$ [ETE $\vee$ ¬ETE].

W13. ETE $\Rightarrow$ J*(H)

From (W13) and CJ:

W14. ETE $\Rightarrow$ (If I introspect, then J*(¬ETE $\vee$ H)).[4]

According to the revised construal of luminosity:

W15. ¬ETE $\Rightarrow$ (If I introspect, then J*(¬ETE)).

From (W15) and CJ:

W16. ¬ETE $\Rightarrow$ (If I introspect, J*(¬ETE $\vee$ H)).

But now the proper conclusion from (W12), (W14) and (W16) is no stronger than:

W17. If I introspect, J*(¬ETE $\vee$ H).

or, maybe:

W18. $\Box$ [If I introspect, J*(¬ETE $\vee$ H)].

Crucially, (W17) and (W18) fall short of (W6) and the conclusion (W7). The key point is that introspection is the source of our (empirically) justified beliefs about what evidence we have. If we bear that in mind, we should be hesitant about accepting Weatherson's argument that I am justified a priori in believing that my total evidence isn't is misleading.[5]

1. Weatherson builds up to this argument over a number of pages. The final variant, which I am following here, is set out at (2005, p. 325-326). One difference between Weatherson's original formulation and my retelling is that, for Weatherson, the initial assumption is "all my possible evidential states are G or not G", where "having evidence that is G justifies belief in snow next winter" (Weatherson's choice of H) (2005, p. 326). Also, he tries to address objections to luminosity assumptions due to Williamson. The luminosity principle Weatherson adopts is: $M \Rightarrow K\neg KK(\neg M)$. (2005, p. 325). The discrepancy between Weatherson's more hedged luminosity assumption and the simpler one I use in the text doesn't affect anything which follows.

2. In the interests of clarity, I will set aside complications that arise from the contingency of my existing and of my having epistemically significant mental states.

3. One might worry that what I am saying here holds for doxastic justification, but not for propositional justification. Then, if we take Weatherson's argument to be about propositional justification, the diagnosis I am presenting will seem to be off the mark. I would dispute this way of seeing things for a number of reasons. However, the issues become complicated quite quickly; I hope that the following observation will suffice. Ultimately the non-Dretskean parties to this discussion will want to say that we can sometimes *know* that our evidence isn't misleading (Weatherson agrees: 2005, p.327). If not, then either the closure principle for knowledge is violated, or skepticism prevails. It seems pretty uncontroversial to me that you wouldn't *know* any facts about your current mental state if you couldn't or didn't introspect. (Note, too, that Weatherson's own version of luminosity is framed as claim about knowledge). So, at the least, the qualification to the luminosity requirement as presented in the text holds for the sort of justification that is necessary for knowledge. That is enough to put (W1)-(W7) into doubt. Two caveats: I'm setting aside facts about your mental states that you might discover by observing your behavior. Also, I'm not taking for granted that introspection is some kind of inner sense. For useful discussion that bears on the issues here see B. Mc Laughlin and Michael Tye, "Is Content-Externalism Compatible with Privileged Access?", Philosophical Review, July 1998.

4. We have assumed that: (i) $ETE \Rightarrow J^*(H)$. By CJ, (ii) $ETE \Rightarrow J^*(ETE \vee H)$. As a matter of logic, (iii) $(ETE \vee H) \Rightarrow$ (If I introspect, then $(ETE \vee H)$). From (ii), (iii) and CJ, $ETE \Rightarrow J^*$(If I introspect, then $J^*(\neg ETE \vee H)$).

5. There may be other problems with the argument as well. I wouldn't be at all surprised if (W16), (W17), and/or (W18) are still too strong. But, even so, the point stands that the argument for (W7) is defective insofar as it makes use of an inflated version of luminosity.