

Penultimate version; final version in *Noûs* 31 (1997). Please feel free to email me for a pdf of the published version: david.christensen@brown.edu.

What is Relative Confirmation?*

1. Introduction

It is commonly acknowledged that, in order to test a theoretical hypothesis, one must, in Duhem's phrase, rely on a "theoretical scaffolding" to connect the hypothesis with something measurable. Hypothesis-confirmation, on this view, becomes a three-place relation: evidence E will confirm hypothesis H only relative to some such scaffolding B. Thus the two leading logical approaches to qualitative confirmation--the hypothetico-deductive (H-D) account and Clark Glymour's bootstrap account--analyze confirmation in relative terms. But this raises questions about the philosophical interpretation of the technical conditions these accounts describe. What does it mean to say that E confirms H "relative to B"? How should we interpret the relation we are trying to analyze?

The answer to this question must lie largely in how "confirmation-relative-to-B" is to inform rational belief. The relationship between relative confirmation and belief-worthiness may not be all there is to relative confirmation. But if an account of relative confirmation is to have any interest, there should be circumstances in which the fact that E confirms H relative to B should provide reason for believing H. Thus one question that must be answered, if we are to understand fully what we are doing in relative confirmation theory, is "when does the fact that E confirms H relative to B constitute reason for belief in ('real confirmation' of) H"?¹

Now it clearly cannot be that one should take E to provide a reason for believing H whenever there is some B or other relative to which E confirms H. It is only if B meets certain conditions that confirmation relative to it will provide reason for belief. Just what those conditions are, however, is not obvious. Moreover, writers developing relative confirmation theories have surprisingly little to say on this question (though their arguments often seem to presuppose one or another answer to the question).

After examining some attractive answers, I will argue that the choice of an interpretation of relative confirmation deserves more careful thought than it has typically been given. It profoundly affects our assessments of the success or failures of particular relative confirmation theories in

accounting for examples of scientific reasoning. Moreover, it will in large measure determine the interest of the whole project of analyzing relative confirmation.

2. The Truth Requirement

What is one assuming when one uses certain background beliefs B in testing a hypothesis? A natural answer is that one is supposing that B is true. Thus one obvious answer to the question of how relative confirmation relates to increased belief-worthiness is:

The Truth Requirement: The fact that E confirms H relative to B provides real confirmation of H just in case B is true.²

Acceptance of something like the Truth Requirement is suggested by criticisms of particular relative confirmation accounts on the grounds that they allow the confirmation of virtually any H relative to some true B. (See Glymour's (1980, 36) criticism of the H-D theory or Thomas Grimes' (1987, 103-104) criticism of bootstrapping.) If confirmation relative to true background was not supposed to yield real confirmation, the force of these criticisms would be quite unclear.

There is something odd about the Truth Requirement, however. The epistemic import of the relative confirmation of H by E is contingent on the metaphysical matter of B's truth--a matter that is entirely independent of whether there is evidence for B, or, indeed, whether B is even believed to be true. Thus consider two scientists with identical evidence living in different worlds. Suppose that the evidence in each world supports B, but that in only one of these worlds is B really true. Each scientist believes B, knows that E confirms H relative to B, and, upon discovering that E is true, increases her confidence in H. According to the Truth Requirement, only one of these scientists was acting in accordance with reasons for belief provided by E's confirming H. The Truth Requirement, then, cannot provide an acceptable account of the relation between relative confirmation and belief-worthiness.

3. The Confirmation Requirement

A better suggestion is that a scientist's good reason to believe B, rather than the truth-value of B, that is relevant. Interpreting "reason to believe" as confirmation yields the following:

The Confirmation Requirement: The fact that E confirms H relative to B provides real confirmation of H just in case (or to the extent that) B is confirmed.

The first thing we need to get clear on here is what is meant by "confirmed." Clearly, we cannot simply mean that B is confirmed by another instance of relativized confirmation. The fact that B is confirmed by some E' relative to some B' or other gets us nowhere. The requirement's plausibility derives from taking the confirmation of B to be the kind which itself licenses belief--i.e., real confirmation.

Moreover, this real confirmation cannot simply be provided by relative confirmation that the Confirmation Requirement has sanctioned. If one saw all real confirmation as arising from relative confirmation validated by real confirmation of B, one would be involved in vicious regress (B's real confirmation would depend on its confirmation relative to a really confirmed B', whose real confirmation would in turn depend on its confirmation relative to a really confirmed B'', etc.). Thus the Confirmation Requirement needs a separate account of real confirmation--something in addition to the real confirmation provided by relative confirmation satisfying the Confirmation Requirement.

Now there is something troubling in the suggestion that relative confirmation plus antecedent real confirmation gets you real confirmation. If we had in hand a satisfactory theory of real confirmation, why would we be interested in relative confirmation at all? However, there may still be interesting relative confirmation projects that would mesh with the Confirmation Requirement.

One might take real confirmation to be achievable in a way dependent on relative confirmation--for instance, by taking certain complicated structures of interlocking relative confirmation to constitute real confirmation of a set of hypotheses. Glymour may have something like this in mind for his bootstrap account of confirmation. Bootstrap confirmation is relative to theories, and although Glymour declines to give a full account of theory-acceptance, his informal criteria for theory-choice all involve bootstrap confirmation of the theory's constituent hypotheses.³

Another approach might recognize a structurally "direct" mode of real confirmation--one that did not depend on relative confirmation at all. Duhem thought that whole systems of hypotheses could be tested in a non-relative way. However, Duhem's method of non-relative confirmation could not be applied to typical theoretical hypotheses singly. Thus on some kinds of accounts, a substantial job would be left for a relativized theory of hypothesis confirmation to perform.

These possibilities suffice to show that there is nothing incoherent about pursuing relative confirmation understood along the lines of the Confirmation Requirement. However, in doing so, one commits oneself in advance to there being some additional way of achieving real confirmation that meets two conditions: first, it must be distinct from the way provided by relative confirmation satisfying the Confirmation Requirement; second, it must leave some work for relative confirmation to do. This liability may not be fatal. But it does provide reason for considering alternatives to the Confirmation Requirement.

4. The Belief Requirement

The difficulties with the Confirmation Requirement flow from its dependence on an additional theory of real confirmation to validate B. A natural suggestion for eliminating the difficulties, then, is to eliminate the demand that B be validated, as follows:

The Belief Requirement: The fact that E confirms H relative to B provides real confirmation of H just in case (or to the extent that) B is believed.⁴

Something like the Belief Requirement seems presupposed by criticisms of bootstrapping on grounds that it excludes confirmations that would be intuitively respectable (or licenses confirmations that would be intuitively bogus) for those who believed B. (See Christensen (1990, 646-647, 651-652).) Obviously, this account does not threaten to land us in vicious regress as does the Confirmation Requirement. And it would seem to capture at least one of the central intuitions about relative confirmation: that the way we should take E to bear on H is determined by the connections we believe to hold between the two.

Of course, the Belief and Confirmation Requirements will offer different assessments in cases where E confirms H relative to B, but where B is believed without being confirmed. And here, it might seem that E shouldn't provide the agent with reason to believe H. If an agent's belief in H is in large measure based on its confirmation relative to some irrationally believed B, few would take the agent's belief in H as beyond rational criticism, even if its confirmation relative to believed B was very strong.

On the other hand, such a belief is intuitively more responsible than it would be if the agent didn't even believe B. Indeed, if H is supported by E relative to background beliefs the agent has accepted, she would be making an epistemic mistake not to give H more credence. This seems independent of whether the agent has made a different epistemic mistake by accepting B in the first place. Thus in cases where E confirms H relative to B, the agent's belief in B may give her some reason for taking E to support H.

Relative confirmation understood according to the Belief Requirement, then, seems to capture an important aspect of confirmation: how what you do believe affects what else you should believe. Relative confirmation interpreted according to the Confirmation Requirement seems to capture an additional aspect: how what you should believe affects what else you should believe. A fully rational belief may not be describable purely in terms of its confirmation relative to believed background. But confirmation relative to what the agent believes is plausibly an important component of rational belief.

Although these two Requirements provide importantly different interpretations relative confirmation, it is worth noting that the intuitive tests we use in evaluating particular relative confirmation theories will often--perhaps usually--be indifferent to which understanding we employ. Most examples we use to test relative confirmation theories involve scientists who have accepted a certain theory. Though it is not usually mentioned, it may usually be assumed that the scientists accepted the theory for good reasons. In such cases, the two interpretations yield the same intuitive test: We consider a real or imagined example of putative confirmation, and ask ourselves "Would a reasonable B-proponent take E as a reason for believing H?" On either the

Belief or the Confirmation Requirement understanding, our theory of relative confirmation should see E as confirming H relative to B in the same cases in which the answer to our intuitive question is "yes." Interestingly, both leading accounts of relative confirmation have run afoul of exactly this sort of intuitive test.

5. Relative Confirmation and Epistemic Interdependence

Suppose that an agent believes background theory B. We may even suppose that B is well-confirmed for him (and, for that matter, true). Even in this situation, B may be intuitively inappropriate for allowing certain bits of evidence to support belief in certain hypotheses, because of epistemic interdependencies between the various terms of the relative confirmation relation.

Let us consider first an example Glymour raised against a simple H-D account of relative confirmation:

The hypothetico-deductive account of confirmation yields ... the following absurdity.

Given a true evidence sentence e , almost every sentence S is confirmed by e relative to some true theory. More precisely, let e be a true not logically valid sentence, and S any sentence at all that is both consistent and not a logical consequence of $\sim e$; then S is confirmed by e with respect to $(S \supset e)$, which is true. (1980, 36)

The absurdity here is not that any sentence can be confirmed relative to some theory or other--no one would think that confirmations relative to arbitrary theories should be taken seriously. Glymour emphasizes that the theory in question is true, suggesting that something like the Truth Requirement is implicit here. But the problem is not dependent on the Truth Requirement. If a logically competent agent reasonably believes e , she will also reasonably believe $(S \supset e)$; thus this relative confirmation should count as real on either the Confirmation or Belief Requirements.

Thinking in these terms highlights the source of the absurdity. The background theory is one the agent would reasonably believe simply because she knew E to be true. But this undermines the basic H-D intuition. In H-D confirmation, E is taken as evidence for H roughly because E is what we would expect if H were true. But in this case, our reason for believing that E will be true

if H is true is simply that know E to be true anyway. Since H plays no role in showing that E was to be expected, no intuitive confirmation takes place. If the agent had reason to believe $(S \supset e)$ independent of her reason for believing e, then the confirmation would be perfectly legitimate. Thus it is the intuitive dependence of B on E that undermines confirmation in Glymour's case.⁵

A similar kind of example was used by Christensen (1983) against the original bootstrap account. On this account, the observation of a black raven (Ra & Ba) turned out to confirm pantheism $(x)Gx$, relative to a theory T whose axioms are the pantheistic principle and the famous raven hypotheses $(x)(Rx \supset Bx)$. Again, the problem is not just that there exists some theory relative to which sighting a black raven confirms pantheism. The problem arises because, intuitively, sighting a black raven typically should not confirm pantheism for a pantheist who also believes that all ravens are black. Even if we assume that pantheism and the raven hypothesis are reasonably believed by the agent (and even if we assume that T is true) it is intuitively wrong for the agent to take a black raven as confirming pantheism. The problem is this: the fact that E confirms H relative to T does not provide real confirmation of H, even for an agent who believes T, even for an agent for whom T is confirmed (and even if T is true). The example shows that the original bootstrap account fails to measure up to the standards set by the Belief or Confirmation (or even Truth) Requirements.

In this case, the bogus confirmation uses as an auxiliary a bizarre consequence of the theory: $(x)[(Rx \supset Bx) \equiv Gx]$. This auxiliary is intuitively inappropriate for confirming pantheism because we assume that the agent's only reason for believing it is that she already subscribes to pantheism. As in the previous case, if she had independent reason to believe in the auxiliary, the confirmation would be legitimate. But a pantheist who also believes that ravens are black will not typically have any such reason. Once again, the intuitive lack of confirmation stems from epistemic dependencies among the terms of the relative confirmation relation.

One lesson we could take from these examples, then, is this: a theory of relative confirmation interpreted according to the Belief or Confirmation Requirements must be able to

make certain discriminations. It must be sensitive to confirmational interdependencies among B, H, and E, in a way that H-Dism and the original bootstrap account are not.

Although it's not clear just how Glymour interprets relative confirmation, he certainly took each of the above examples to refute the confirmation theory that licensed it. He and others (see Glymour (1983a); Zytlow (1986); Earman and Glymour (1988)) then reformulated bootstrapping to eliminate the kind of circularity we see in the ravens-pitheism example. However, Christensen (1990) showed that the reformulations allowed essentially similar circular confirmations. None of these authors were very explicit about how relative confirmation was to be interpreted. Nevertheless, the intuitive tests they employed seem to presuppose that a satisfactory confirmation theory should not allow relative confirmation in cases where reasonable proponents of the background theory would not see real confirmation. Thus something like the Belief or Confirmation Requirement seems to underlie these evaluations of bootstrapping's success as a theory of relative confirmation.

But this is not the only possible response to the examples. Another is to reject the idea that relative confirmation should be interpreted according to the Requirements we've examined so far. A different interpretation might yield a dramatically different assessment of putative counterexamples. Let us now examine this alternative.

6. The Independent Confirmation Requirement

Paul Horwich (1983) argues that Glymour's above-quoted counterexample to H-D is innocuous:

For if S is an arbitrary truth, then the so-called background theory, $(S \supset e)$, cannot have been independently established. Therefore this relative confirmation of S should not enhance its credibility. (58)

Horwich argues that it doesn't count against H-Dism that it allows relative confirmation in this case where we see no real confirmation. Yet the background in this instance of relative confirmation is well-confirmed. This is clearly incompatible with interpreting H-Dism according to the Confirmation Requirement. (The same would hold for the Belief and Truth Requirements.)

Instead, the passage suggests an alternative understanding of relative confirmation: that it should yield real confirmation only when B is confirmed independently of E.

A more developed interpretation along these lines is suggested in Edidin (1988). Edidin also dismisses Glymour's example on the grounds that the auxiliary used in the confirmation is not "credible independently of the evidence invoked" (268). He goes on to dismiss the Christensen (1983) counterexample to bootstrapping on the grounds that its auxiliary does not have "antecedent credibility independent of that of the hypothesis" (269). Edidin writes:

Neither Glymour's purported counterexample to the H-D account of relative confirmation nor Christensen's to Glymour's own bootstrapping account succeeds in discrediting its target. The reason is the same in both cases: failures of confirmation need not be failures of relative confirmation. In each instance, it seems clear that any adequate account of how relative confirmation yields confirmation simpliciter will rule out the undesirable cases even if they are genuine instances of relative confirmation. (270)

Thus the independence conditions violated in the counterexamples must, according to Edidin, be built into our account of when relative confirmation yields real confirmation.

Most recently, Sam Mitchell (1995) writes:

Glymour can reply to every counter example so far proposed if he adopts the following principle:

No graph confirms its hypothesis absolutely if it is unlikely that any reply can be given to [a challenge to the auxiliaries used in the confirmation graph] by independently supporting auxiliaries. (244)

The following seems to incorporate the spirit of the suggestions made by Horwich, Edidin, and Mitchell:

The Independent Confirmation Requirement: The fact that E confirms H relative to B provides real confirmation of H just in case (or to the extent that) B (or any hypothesis from B used as an auxiliary) is confirmed independently of E and H.⁶

This interpretation of relative confirmation clearly merits attention. After all, it seems to provide a direct and simple answer to troublesome counterexamples. But before moving on to a closer examination of its philosophical merits, I want to bring out another reason for taking the interpretation seriously. In the next section, I'll argue that a very different approach to handling the counterexamples--an approach which seems at first to have no special connection to the interpretive issues we have been concerned with--actually turns out to require understanding the relative confirmation relation in Independent Confirmation Requirement terms.

7. The Subtraction Strategy

Madison Culler (1995) and Ken Gemes (1994) offer accounts of relative confirmation developed from bootstrapping with an eye to avoiding the counterexamples. Neither explicitly endorses a particular account of how relative confirmation relates to real confirmation. However, both use the same prima facie attractive strategy for dealing with the counterexamples.

The counterexamples involve confirmation relative to a theory T which includes the hypothesis H to be tested--a possibility (dubbed "macho bootstrapping" in Earman and Glymour (1988)) that Glymour originally wanted to allow for, but has since rejected. Both Culler and Gemes reject macho bootstrapping, and offer accounts of relative confirmation which are designed to apply only to cases where T does not entail H.

In conformance with this strategy, Culler and Gemes show how the problematic examples could be handled, once purged of H's entailing T. This involves a two-step process. First, the examples are altered by subtracting H from T. Confirmation is then analyzed relative to the reduced background theory T'. Let us call this "the subtraction strategy" for dealing with the problematic examples.

Of course, there is more than one way to subtract H from T: various different sets of sentences T', when conjoined with H, would yield the original T. Both Culler and Gemes subtract H from T in the intuitively natural way. The natural way removes not only H, but also those elements--including the problematic auxiliaries--whose plausibility intuitively depends on H. Removing these elements disallows relative confirmation in the problem cases.

But does this solve the difficulty posed by the counterexamples? Suppose that one subtracted H in an intuitively unnatural way. For example one might subtract H from T in the raven/pantheism example by letting T' be just the bizarre auxiliary. Both Culler's and Gemes' accounts would then allow confirmation. But given that our pantheist believer in the raven hypothesis accepts the auxiliary only because he accepts pantheism, such confirmation is just as intuitively bogus as confirmation relative to the unreduced T.

Does this sort of example then show Culler's and Gemes' accounts to be unacceptable? The answer depends on the question that we have been focussing on: how is confirmation relative to T' to be interpreted? When should such relative confirmations count as real? In order to avoid sanctioning the bogus confirmations as real, Culler and Gemes would need to hold that confirmations of H relative to T' provide real confirmation only if T' is independent of H.

Holding this is, of course, inconsistent with interpreting relative confirmation according to the Truth, Belief, or Confirmation requirements. T' may well be true, believed, and confirmed even when parts of it are dependent on H. Thus it turns out that one very attractive strategy for responding to the counterexamples--eschewing macho bootstrapping, and subtracting H from T before assessing relative confirmation--succeeds only given a particular understanding of relative confirmation. It requires interpreting relative confirmation in Independent Confirmation Requirement terms.

To put the point more generally: even if macho relative confirmation is never legitimate, we cannot solve the problem disclosed in the counterexamples simply by prohibiting T from entailing H. For there will still be plenty of cases where T does not entail H, but where our reasons for believing (relevant parts of) T are that they follow from H. Unless our account of relative confirmation can catch such intuitive dependencies, we will need to embrace the Independent Confirmation Requirement if we are to avoid claiming real confirmation in these cases.

8. Can the Other Requirements Simply be Dismissed?

Should we, then, simply adopt the Independent Confirmation Requirement interpretation of relative confirmation? There is little discussion of this question in the literature. But some have made claims which, if accepted, would rule out at a stroke all of the competitors we've encountered so far.

Let us begin with a passage from John Earman's recent Bayes or Bust. Earman is commenting on the significance of some counterexamples to bootstrapping:

... we need to know what end the three-place Glymourian relation 'E bootstrap-confirms/tests H relative to T' is to be put. E bootstrap-confirms H relative to T cannot be taken to imply that, assuming T to be true or well-confirmed, E confirms H, for in the cases at issue H is part of T. Rather, the most plausible usage is in adjudicating questions of evidential relevance. (1992, 75)

Earman's argument would reject the Truth and Confirmation Requirements, at least as applying to any account that allowed macho relative confirmation. (Presumably, the argument would apply to the Belief Requirement as well.)

However, I do not think we should find argument persuasive. It seems to assume that if H is taken to be well-confirmed (or even true), the question of whether E really confirms H is somehow rendered moot. But why should this be so? Confirmation theory is not just supposed to help us decide whether the evidence tells in favor of hypotheses about which we are undecided. It is also supposed to tell us which of our currently accepted hypotheses are confirmed by which particular bits of evidence.⁷

Another claim that would rule out the Independent Confirmation Requirement's competitors is made by Edidin in the passage quoted above: "any adequate account of how relative confirmation yields confirmation simpliciter will rule out the undesirable cases." The undesirable cases in question are those in which confirmational interdependencies between H and T or E undermine the legitimacy of confirmation. If Edidin's claim is true, then the Truth, Confirmation, and Belief Requirements are simply non-starters, for they make no attempt to weed out the problematic interdependencies.

But to agree--as everyone presumably does--that the problematic confirmations must be weeded out somehow is not to agree that they must be weeded out by the account of how relative confirmation yields real confirmation. There is another possibility: they may be weeded out by the account of relative confirmation itself. Glymour's account, for instance, is designed to do just that. Edidin presents no explicit argument for the strong claim that the problematic examples must be weeded out by the account of how relative confirmation yields real confirmation. But one possible reason for holding it is suggested in the following passage:

[A]n account of relative confirmation cannot by itself determine the conditions under which a hypothesis is genuinely confirmed. This limitation, though it restricts the interest of accounts of relative confirmation alone, also serves to shield such accounts from certain potential counterexamples. (1988, 266)

Perhaps the idea is this: theories of relative confirmation are, by their very nature, incapable of being sensitive to the problematic interdependencies. This claim is further suggested by Edidin's comment that when Glymour criticized the H-D account, and when others criticized bootstrapping they were "guilty of failing to appreciate the significance of the gap between confirmation relative to auxiliaries and confirmation simpliciter". (1988, 268)

But why should we believe this claim? As noted above, revisions of bootstrapping did eliminate the raven/pantheism example. True, these revisions failed to block similar counterexamples, but for all we have seen, other emendations, also within the relative confirmation framework, might yet succeed. After all, the confirmation theories in question all presuppose that confirmational relations are largely based on logical relations. If this is true, it is hard to discount a priori the possibility that some syntactic test within a relative confirmation theory will be able to disclose the problematic confirmational interdependencies.⁸ Thus it seems to me that a proponent of the Independent Confirmation Requirement may not dismiss its competitors out of hand.

9. Assessing the Independent Confirmation Requirement

Let us begin by examining Glymour's discussion of a proposal for fixing the H-D model (1980, 38-39). The proposal is that, for E to confirm H relative to T, we should require that H be

necessary to deduce E. In other words, we would require that the evidence not be derived from that portion of T that would be left over when H was subtracted from T. Thus the proposal involves something closely akin to the subtraction strategy.

In rejecting the proposal, Glymour notes that there are many ways to cut up a theory. Whether H is necessary to deduce E depends on what one takes the remainder T' to be. Now some ways of cutting T up would be unnatural, but, Glymour notes, H-D theorists had no account of naturalness. He takes this as undermining the proposal's viability.

But why should the lack of an account of naturalness be such a big problem? Glymour doesn't elaborate, but the answer, I take it, is this: Such a move would involve helping oneself to a large portion of what one had set out to explain. What makes one way of cutting up a theory natural is largely determined by what parts of the theory depend on what other parts. In other words, it is the confirmational interdependencies among propositions that account for our judgements of naturalness in these matters. But it was precisely such confirmational interdependencies that the H-D account was supposed to explain!

Precisely analogous problems arise for the uses of the subtraction strategy considered above. In deciding whether to take an instance of relative confirmation seriously, we must rely on judgements about which parts of T are plausible independent of H--our judgements about their confirmational interdependencies. But these interdependencies are part of the intended explicandum for confirmation theory. Thus the advantages secured by use of the subtraction strategy are purchased at a price--the price of reducing the scope of relative confirmation theory significantly.

This point applies to any account which embraces the Independent Confirmation Requirement. Before knowing whether we are to take E's confirmation of H relative to B as a reason for believing H, we must first know whether our reasons for believing B are independent of our reasons for believing E and H. Again, this was just the sort of information that relative confirmation theory was supposed to provide. Thus the importance of getting clear about one's

interpretation of relative confirmation goes well beyond simply avoiding confusion or escaping counterexamples.

We saw above that adopting the Confirmation Requirement commits one to the existence of an additional account of real confirmation, and that this account of real confirmation has to satisfy two conditions: First, to avoid regress, it must be independent of the sort of real confirmation provided by relative-confirmation-satisfying-the-Confirmation-Requirement. Second, it must leave significant work for the relative confirmation theory to do.

The Independent Confirmation Requirement commits one to the existence of an additional account of real confirmation meeting even more stringent conditions. To employ the Independent Confirmation Requirement in a given case, we must first determine whether B is really confirmed independently of E and H. Thus our additional account of real confirmation must not only determine real confirmation of B, it must discern the real confirmational interdependencies between E and H and B (or particular hypotheses in B). This is a considerably taller order; the Independent Confirmation Requirement presupposes the existence of a particularly detailed and powerful additional account of real confirmation.⁹

The need for such powerful supplementation threatens the philosophical interest of the relative confirmation project. For suppose we had in hand a successful account of the real confirmational interdependencies among specific hypotheses and evidence statements. Would we need to augment it by an account of relative confirmation understood according to the Independent Confirmation Requirement? The point of such an account would presumably be to help us understand the (real) confirmational status of particular hypotheses. But if we are starting out with an independent account that already captures the real confirmational interdependencies among specific hypotheses and statements of evidence, then it is not clear that our account of real confirmation via relative confirmation sanctioned by the Independent Confirmation Requirement is left with any interesting work to do.

Most of those who advocate the Independent Confirmation Requirement, or whose accounts turn out to require it, do not offer any account at all of how their relative confirmation theories

could be supplemented by the sort of additional account of real confirmation we have seen them to require. The one exception is Mitchell, in his above-mentioned version of bootstrapping. Mitchell suggests, roughly, that real confirmation of auxiliaries occurs when they are supported by a loopless tree of further auxiliaries, beginning with auxiliaries confirmed relative to only themselves. (Glymour's account allows in certain circumstances for a hypothesis to figure in its own confirmation.)

If this suggestion worked, it would provide a nice, simple answer to the troubling questions raised by the Independent Confirmation Requirement. However, we have so far seen no reason to believe that the trees Mitchell's account requires actually exist in cases where we see legitimate confirmation. Mitchell cites examples of reasonable self-relative confirmations in Glymour (1980) and Zytzkow (1986). But these examples do not exhibit the kind of theoretical structure present in the central sorts of cases bootstrapping was designed to cover. Indeed, it is hard to see how the requisite sorts of trees could be constructed for even the artificially simple kinds of examples the discussion has focussed on so far. Moreover, even if the trees Mitchell describes do exist in typical cases of confirmation, the success of the proposal would also require that such trees not exist in cases where we do not see legitimate confirmation. Thus, at this point, although Mitchell's loopless tree account has the form of an additional theory of real confirmation appropriate for use with the Independent Confirmation Requirement, we have no evidence that the account succeeds in discriminating between good and bad scientific reasoning.¹⁰

In sum, then: although pursuing relative confirmation theory under the Independent Confirmation Requirement seems quite attractive at first, it places a very heavy burden on one's account of the conditions under which relative confirmation yields reasons for belief. What's more, if the burden is met, this may well render one's relative confirmation theory superfluous. And at this point, no author has given us much reason to believe that the sort of additional account of real confirmation required by the Independent Confirmation Requirement exists.

10. Conclusion

The Belief, Confirmation, and Independent Confirmation Requirements may be seen as falling on a kind of continuum. As we move along the continuum, we put more and more demands on our account of how relative confirmation yields real confirmation. As this happens, of course, the job assigned the relative confirmation theory becomes easier to perform--and, proportionately, less interesting. Of course it would be better to have an account which met the Belief or Confirmation Requirements than one which met the Independent Confirmation Requirement--such an account would provide a much more complete account of the phenomenon we seek to explain.

At this point, we do not know whether a logical account of relative confirmation meeting the tougher standards is possible. Such an account will have to include a solution to the problem posed by confirmational interdependencies among hypotheses, evidence, and background assumptions--a difficult undertaking. But the program of giving a logical account of confirmation built around a relative confirmation notion satisfying only the Independent Confirmation Requirement cannot avoid this difficult undertaking. For the relative confirmation notion involved in such a program must be supplemented by an account of real confirmation which itself solves the problem posed by confirmational interdependencies. At this point, it is not clear which research program in logical confirmation theory--if either--will ultimately bear fruit.

But if no unique answer to our title question has been established, the importance of thinking clearly and explicitly about the question has. We should understand how our choice of interpretation--or our choice of a confirmation-theoretic strategy that requires a certain interpretation--may serve to affect more than the vulnerability of our account to certain counterexamples. We must be wary of implicitly helping ourselves to what we are trying to explain, and we must be careful to distinguish between solving the problem of accounting for confirmational interdependencies, and acting as if the problem has been solved already. The interest of any theory of relative confirmation is intimately bound up with how it understands the claim that E confirms H "relative to B." To choose one's interpretation of relative confirmation is, in large measure, to set one's philosophical sights.

NOTES

* I would like to thank Ken Gemes, Steve Jacobson, Hilary Kornblith, Sam Mitchell, Derk Pereboom, and an anonymous referee for Noûs, for very helpful comments on earlier drafts. I'd also like to thank the audience at Mount Holyoke College, where an earlier version of this paper was presented. And I'd like to thank the University of Vermont for research support.

¹ The importance of understanding the relation between relative confirmation and credibility is nicely emphasized in Edidin (1988). Edidin's views on the subject will be discussed below.

² Strictly speaking, this principle should be understood as relativized to an agent who knows E. For brevity, I'll omit mention of this when stating principles relating relative confirmation to real confirmation.

³ See Glymour (1980) 152-155. Glymour does not specify which theory these bootstrap tests are to be relative to. My guess is that it is intended that theory-choice is ultimately to be grounded in testing the hypotheses in a theory relative to the theory itself.

⁴ In addition to the relativization noted in fn. 2 above, the Belief Requirement should be understood as relativized to the agent who believes B.

⁵ This diagnosis follows Edidin (1988).

⁶ The second parenthetical phrase corresponds to Edidin's comments on the counterexample to bootstrapping. Bootstrap confirmation is defined relative to theories. But Edidin's complaint about the example was not that the theory involved in lacked independent credibility; it was that the auxiliary from the theory used in the problematic confirmation was not independently credible.

⁷ In the quote's last line, "adjudicating questions of evidential relevance" seems to be offered as a use for relativized confirmation theories which is an alternative to using them in analyzing real confirmation. But evidential relevance is a very slight generalization of confirmation: E is relevant to H iff E confirms or disconfirms H. Thus evidential relevance can be thought of in relativized or non-relativized form. A relative account of confirmation will trivially yield a relative account of evidential relevance; surely this cannot be the ultimate purpose to which our relative account of confirmation is put. But in order to yield a non-relativized account of evidential relevance, a

relative confirmation theory will have to yield a non-relativized account of confirmation already. Thus using an account of relative confirmation to adjudicate questions of evidential relevance requires, and cannot be an alternative to, giving an account of how relative confirmation yields real confirmation.

⁸ One possibility along these lines is to abandon the classical approach to theories. Glymour, for example, represents a theory by the deductive closure of the natural first-order representations of its intuitive axioms. If we jettisoned this approach, we might be able to avoid seeing bizarre auxiliaries like $(\forall x)((Rx \supset Bx) \equiv Gx)$ as on a par with a theory's basic principles, and might thereby avoid the problematic confirmations. (The possibility of avoiding the counterexamples in this way is discussed in Christensen (1983, 1990)).

One example of such an approach would be to take the closure of the set of basic principles in some version of relevance logic. For an argument that relevance logic could be used to solve evidential relevance problems in H-Dism, see Waters (1987). The suggestion that relevance logic could be used to solve bootstrapping's problems has also reportedly been made by Kevin Kelly; see Glymour (1983b).

Gemes' version of H-Dism (1993) avoids the Glymourian counterexample without invoking the subtraction strategy. Thus it, unlike his version of bootstrapping, seems amenable to interpretation along the lines of the belief or confirmation requirements.

⁹ In section 3, we saw two ways a theory of real confirmation for B might be added to an account of relative confirmation without incurring regress difficulties: taking certain interlocking structures of relative confirmation to constitute real confirmation of a set of hypotheses; and taking certain large sets of hypotheses to be confirmable in a direct way. It is doubtful that either of these ways of really confirming whole sets of hypotheses will yield the judgements about real confirmational interdependencies among specific hypotheses (and evidence statements) that the Independent Confirmation Requirement depends on.

¹⁰ I should note that Mitchell himself expresses doubts that his account will provide a purely formal solution for bootstrapping's problems.

REFERENCES

- Christensen, D. (1983) "Glymour on Evidential Relevance," Philosophy of Science 50: 471-481.
- (1990) "The Irrelevance of Bootstrapping," Philosophy of Science 57: 644-662.
- Culler, M. (1995) "Beyond Bootstrapping: A New Account of Evidential Relevance," Philosophy of Science 62: 561-580.
- Earman, J. (1992) Bayes or Bust?: a Critical Examination of Bayesian Confirmation Theory (Cambridge, MIT Press).
- and C. Glymour (1988), "What Revisions does Bootstrap Testing Need?," Philosophy of Science 55: 260-264.
- Eddin, A. (1988) "Discussion: From Relative Confirmation to Real Confirmation," Philosophy of Science 55: 265-271.
- Gemes (1993) "Hypothetico-Deductivism, Content, and the Natural Axiomatization of Theories," Philosophy of Science 60: 477-487.
- (1994) "Bootstrapping and Content Parts," unpublished manuscript.
- Glymour, C. (1980) Theory and Evidence (Princeton: Princeton: Princeton University Press).
- (1983a) "Revisions of Bootstrap Testing," Philosophy of Science 50: 626-629.
- (1983b) "On Testing and Evidence," in J. Earman, ed. Testing Scientific Theories, Minnesota Studies in the Philosophy of Science vol. X (Minneapolis: Univ. of MN Press).
- Grimes, T. R. (1987) "The Promiscuity of Bootstrapping," Philosophical Studies 51: 101-107.
- Horwich, P. (1983) "Explanations of Irrelevance," in J. Earman, ed. Testing Scientific Theories, Minnesota Studies in the Philosophy of Science vol. X (Minneapolis: Univ. of MN Press).
- Mitchell, S. (1995) "Towards a Defensible Bootstrapping," Philosophy of Science 62: 241-260.
- Waters, K. (1987) "Relevance Logic Brings Hope to Hypothetico-Deductivism," Philosophy of Science 54: 453-464.
- Zytkow, J. (1986) "What Revisions does Bootstrap Testing Need?," Philosophy of Science 53: 101-109.