Toward a Synthesis of Reliabilism and Evidentialism?
Or: Evidentialism's Troubles, Reliabilism's Rescue Package

Alvin I. Goldman
Rutgers University

For most of their respective existences, reliabilism and evidentialism (that is, process reliabilism and mentalist evidentialism) have been rivals. They are generally viewed as incompatible, even antithetical, theories of justification.[1] But a few people are beginning to re-think this notion. Perhaps an ideal theory would be a hybrid of the two, combining the best elements of each theory. Juan Comesana (forthcoming) takes this point of view and constructs a position called "Evidentialist Reliabilism." He tries to show how each theory can profit by borrowing elements from the other. Comesana concentrates on reliabilism's problems and how it might be improved by infusions from evidentialism. This paper follows a similar tack. My emphasis, however, is the reverse of Comesana's. I highlight problems for evidentialism and show how it could benefit by incorporating reliabilist themes. I am not sanguine that evidentialists will see it my way. They might even view my proposals as an insidious attempt to convert evidentialists to reliabilism. Well, I won't debate the best way to formulate this paper's recipe. At any rate, it began with the idea (which anteceded my reading of Comesana) of creating a synthesis of reliabilism and evidentialism. It retains significant strands of that idea, although the synthesis theme does not pervade the entire paper.

What is mentalist evidentialism? Its original formulation was succinct:

[EJ] Doxastic attitude D toward proposition p is epistemically justified for S at t if and only if having D toward p fits the evidence S has at t. (Feldman & Conee, 1985/ 2004: 83).

I begin by raising worries about the notions of *evidence* and *fittingness* that are crucial to evidentialism as formulated by [EJ]. To help with many of the problems raised, I recommend supplements borrowed from reliabilism. As we proceed I shall also present reasons why reliabilism might be improved by doses of evidentialism.

1. The Concept of Evidence in Evidentialism

A theory of justification that calls itself "evidentialism" and makes evidence possession its focus can reasonably be asked to explain *which* concept of evidence it means to invoke. If the only viable concept of evidence available to it introduces elements from a "foreign" approach such as reliabilism, this should provide some motivation toward a synthesis or unification with that approach. This is the situation facing evidentialism, as I argue in this section.

In an overview of possible conceptions of evidence, Thomas Kelly (2006) offers several conceptions worthy of attention. First, evidence may be understood as "that which justifies belief." As Jaegwon Kim puts it, "When we talk of 'evidence' in an epistemological sense we are talking about justification: one thing is 'evidence' for another just in case the first tends to enhance the reasonableness or justification of the second" (1988: 390-391). Can this be the sense of the term Feldman and Conee have in mind? Not if their evidentialist theory is intended to provide an analysis or explanatory account of justification. Yet, clearly, this *is* what their evidentialist theory, as expressed in [EJ], is intended to provide. If 'evidence' is defined as "that which justifies belief," then the definition of 'justified' in terms of 'evidence', as proposed in [EJ], is circular and unhelpful.

Next consider a conception of evidence proposed by Timothy Williamson (2000), in which knowledge is equivalent to evidence. At any rate, according to Williamson, all items of evidence are pieces of knowledge and vice-versa. Could Feldman and Conee avail themselves of the sense of 'evidence' in which evidence is equivalent to knowledge? This poses the same problem of circularity as the first definition. Feldman and Conee hold that 'knowledge' is to be analyzed in terms of justification, which, of course, is analyzed in terms of evidence.[2] In light of these commitments, it would again be circular to analyze or define 'evidence' in terms of 'knowledge'.

Another conception of evidence Kelly considers is evidence "as a guide to truth." In other words, something is evidence for p just in case it is a reliable sign, symptom, or mark of the truth of p. This is an appealing conception of evidence, which makes sense of the term's use in many walks of life. In criminal law, it is plausible that what a court

admits as evidence should be items that (when properly interpreted) are reliable guides to truth, or signs of that for which they are taken to be evidence. Similarly, in science the reading of a gauge or instrument is evidence for a certain object's having property F just in case the reading is a reliable indicator of the object's possessing F.

Is this definition one that Feldman and Conee could adopt? It is certainly compatible with certain strands of their approach. There is no looming circularity of the sort that besets the first two definitions. Furthermore, it is compatible with Feldman and Conee's mentalism about evidence, because mental states, events, and processes often qualify as indicators of the truths for which they are (ostensible) evidence. At least this is so if we don't inhabit an evil-demon-world and are not otherwise badly deceived about the actual world. When it visually appears to someone that there is a computer monitor before him, this visual state is usually a reliable sign or indicator that a computer monitor is before him. When a person has an ostensible memory impression that she ate oatmeal for breakfast this morning, the memory impression usually indicates that she did eat oatmeal this morning. Under the truth-indicator conception of evidence, it looks like these two experiences will constitute (prima facie) evidence for just the propositions one would expect them to be evidence for – and that Feldman and Conee take them to be evidence for. So far, then, the reliable indicator interpretation of evidence is one that Feldman and Conee should find congenial.

Feldman and Conee would probably resist this suggestion, of course, for two closely related reasons. First, they hold that even in an evil-demon world a visual experience as of a computer monitor is evidence for the presence of a computer monitor (Feldman, 1985), and one is justified in believing that a computer monitor is present even though such visual experiences are not reliable indicators of the presence of a monitor in such a world. A related problem is that truth-indicatorship properties vary from world to world, so that what is evidence for what is contingent under this approach. But Feldman and Conee want evidential relationships to be necessary.

It should not be assumed, however, that truth-indicatorship in the world of an example is the optimal interpretation of the reliable indicatorship approach. Instead, truth-indicatorship relationships might be *rigidified* so that they are fixed by the correlations that obtain in the actual world. If this approach is adopted, even a cognizer

in an evil-demon world will be justified in his perceptual beliefs (see Goldman, 2008a). Comesana embeds roughly this approach in a two-dimensional semantics (Comesana, 2002; forthcoming) and calls it "indexical reliabilism." In particular, following Stalnaker's (1978) version of two-dimensional semantics, there are two different propositions associated with any attribution of justification: the "diagonal proposition", which implies that the belief is produced by a method that is reliable in the world where it the belief is held; and the "horizontal" proposition, which implies that the belief is produced by a process that is reliable in whichever world the proposition is considered. The latter interpretation might be acceptable to evidentialists, both in terms of allowing perception-based justifiedness in an evil-demon world and in terms of complying with the necessity constraint on the evidential relation.

In this section we have noted that there are multiple conceptions, or definitions, of "evidence." Thus, evidentialism owes us an indication of what *it* means by "evidence." Until we are told how to interpret the term 'evidence' for purposes of this theory, we cannot begin to assess its (extensional) adequacy. As we have seen, moreover, some definitions of 'evidence' are not viable for evidentialism because of circularity. The only definition I am able to pinpoint that does not suffer from this liability employs some sort of reliability notion. Could evidentialism adopt this definition? Pure evidentialism, understood as a rival or opponent of reliabilism, could not adopt it. But when we consider the prospects of a hybrid approach, this might be an attractive option. Mental states would serve as the items of evidence under the hybrid theory, but they would qualify as items of evidence (ultimately) because they -- or their ilk -- stand in reliable-indicator relationships to facts in the world. Although Feldman and Conee themselves are unlikely to applaud this maneuver, other epistemologists in search of the best overall theory might find it a congenial solution to the definitional problem.

2. Justification and Fittingness with Non-Doxastic Evidence

Two notions of justifiedness are commonly recognized in the literature: propositional and doxastic justifiedness. Feldman and Conee's (1985) label for the latter is "well-foundedness."[3] Let us concentrate on this second notion. Feldman and Conee's account of doxastic justifiedness (well-foundedness) is formulated as follows:

[WF] S's doxastic attitude D at t toward proposition p is well-founded if and only if

(i) having D toward p is justified for S at t; and

(ii) S has D toward p on the basis of some body of evidence e, such that

(a) S has e as evidence at t;

(b) having D toward p fits e; and

(c) there is no more inclusive body of evidence e' had by S at t such that having D toward p does not fit e'. (1985/2004: 93)

Obviously, mentalist evidentialism owes us an account of fittingness. Two types of cases need to be covered, fittingness as applied to inferential justifiedness and fittingness as applied to non-inferential justifiedness. I begin with non-inferential justifiedness. Feldman and Conee certainly appear to believe in non-inferential justifiedness. Feldman's textbook *Epistemology* (2003) seems to favor (modest) foundationalism as the best approach to justification, and foundationalism is committed to basic justifiedness, which is non-inferential. Conee's paper "The Basic Nature of Epistemic Justification" (1988/2004), on the other hand, appears to favor a combination of foundationalism and coherentism. Conee agrees, however, that such a theory must acknowledge foundational *experiences* (1988/2004: 43). Thus, both theories seem to be committed to some strand of non-doxastic justifiedness, and, given evidentialism, this must arise from some non-doxastic species of fittingness.[4] Non-doxastic fittingness is what I examine in the present section.

Let us begin with non-inferential beliefs based on perceptual experience or introspection. Almost all foundationalists will accept introspectively-based beliefs (at least some of them) as non-inferentially justified. Modest foundationalists, such as Feldman and Conee, also accept some perceptual beliefs as non-inferentially justified. The question is how, in detail, evidentialism can handle these cases of justifiedness. In both cases the problem I have in mind may be called the *selection* problem. Some perceptually based beliefs are justified and other beliefs are unjustified; similarly for introspectively based beliefs. How can mentalist evidentialism explain which ones are which?

Feldman (2003) is fully aware of the problem and provides a good example. Maurice enters a room, sees a 12-year-old table, and forms two beliefs about it: a belief that it is a table (B(T)) and a belief that it's a 12-year-old table (B(TYOT)) -- although there's nothing in its appearance that hints at its age. Presumably, B(T) is well-founded and B(TYOT) is ill-founded. Why? Feldman wants to say, of course, that the former belief but not the latter "fits" Maurice's experience. But how can fittingness be spelled out so as to yield this result? Feldman tries to explain this fittingness by appeal to the notion of a "proper response to experience." Unfortunately, "proper response" seems like a mere paraphrase of "fittingness," so there is a clear threat of the explanation being unilluminating. Feldman does try to explain in non-circular terms what is meant by a "proper response" to experience. The explanation, however, does not cover all of the relevant territory (see Goldman, 2008b for details). For example, Feldman says that when the contents of a belief are "closer" to the direct contents of the experience, they are more apt to be properly based on experience. This approach assumes that proper responsiveness is always a matter of *content matching*, but this cannot be right. If I form a belief that *this* mental state I am internally demonstrating is a desire, the belief may well be justified. But the belief's content -- that *this* mental event is a desire -- need not match the desire's content, which could be anything, and, in particular, need not concern a desire.[5]

Another move Feldman makes is to appeal to the subject's *training*. He uses this to explain the difference in well-foundedness between an expert and a novice bird-watcher, who both identify a certain bird as of type X. The expert, says Feldman, should be credited with a well-founded belief, the novice (who was just guessing) with an ill-founded one. But how, exactly, is their respective training relevant? A simple and direct answer comes from process reliabilism. The relevance of training, quite simply, is that it makes the bird expert's bird-spotting judgments consistently come out true; the absence of relevant training makes the novice's bird-spotting judgments come out only randomly, or occasionally, true. This is a thinly disguised way of saying that the expert has acquired a reliable process for spotting birds (at least birds of type X); the novice has not acquired any such reliable process, although he lucked out with his present guess. By my lights, evidentialism would be greatly improved if it availed itself of this ingredient of process

reliabilism. Its capacity to illuminate cases of non-inferential well-foundedness could thereby be much enhanced.

The foregoing diagnosis and recommendation also apply to the problem of *introspective* justifiedness. A well-known problem in this arena is Chisholm's (1942) problem of the speckled hen. Exactly which types of belief about one's own present experience are justified? Even if there is a determinate number of apparent speckles associated with a given visual experience, a subject who believes that the hen appears to have exactly 43 speckles will not necessarily be justified in so believing. That depends on how the subject deploys his introspective process and associated processes. Psychologists give the label "subitizing" to a process of simply scanning a set of objects and coming up with a numerical estimate of them. For a small number of objects, subitizing is quite reliable. There is no need for deliberate counting. For larger sets of objects, subitizing is not accurate. Now suppose that a subject uses subitizing to arrive at a belief that his current visual experience features exactly 43 apparent speckles; then he isn't, intuitively, justified -- even if, by chance, 43 is the correct number. This diagnosis can be accommodated within a theory of justification that highlights mental *processes* and their *reliability,* but -- as far as I can see -- only within such a theory. Evidentialism as currently formulated is not such a theory.

I do not mean to reject categorically the notion of fittingness for purposes of a theory of justification. Indeed, Feldman's appeal to content matching in the account of fittingness may hold some merit, at least as part of the story. I return to these topics in a more ecumenical spirit in sections 4 and 7 below.

3. Memory Belief and Current Evidence

Another domain of non-inferential justifiedness to which foundationalists commonly appeal is memory-based beliefs. A familiar idea is that beliefs about the past can be justified by ostensible memories, viewed as non-doxastic states. If I seem to recall eating a bagel for breakfast this morning, this ostensible memory creates justification for my belief that I did eat a bagel this morning. Presumably, Feldman and Conee will be happy to say that memory beliefs of this sort are justified insofar as they "fit" the non-

doxastic evidence provided by a conscious memory experience of this kind. I will not challenge this theory when it comes to beliefs about the past based on such experiences. But another kind of memory – so-called "preservative" memory – has a rather different epistemic role to play in our cognitive life. Preservative memory does not create or generate justifiedness "from scratch," but instead *transmits* a belief's justifiedness (or unjustifiedness) from one time to a later time. Here, it seems to me, evidentialism runs into serious trouble when asked to account for this epistemic role.

Here is an example. Years ago Ichabod formed a belief in proposition Q by acquiring it in an entirely justified fashion. He had excellent evidence for believing it at that time (whether it was inferential or non-inferential evidence). After ten years pass, however, Ichabod has forgotten all of this evidence and not acquired any new evidence, either favorable or unfavorable. However, he continues to believe Q strongly. Whenever he thinks about Q, he (mentally) affirms its truth without hesitation. At noon today Ichabod's belief in Q is still present, stored in his mind, although he is not actively thinking about it. I stipulate that none of his other beliefs confers adequate evidence either for believing Q or for disbelieving it. Since Ichabod remembers Q's being the case, and since he originally had excellent evidence for Q, which was never subsequently undermined, Ichabod's belief in Q at noon today is justified.[6] Moreover, as I argue elsewhere (Goldman, 2009), if we refuse to grant justifiedness to beliefs of this sort, which derive from preservative memory, there will be serious skeptical ramifications: people will fail to know a great many things that common sense credits them with knowing. But does Ichabod's noontime belief in Q satisfy Feldman and Conee's account of well-foundedness? Does it qualify as doxastically justified under their theory?

No. At noontime Ichabod undergoes no evidential *experience* of seeming to remember Q and (by hypothesis) possesses no stored beliefs on which his belief in Q is based -- or which it fits. Years ago, as indicated, he had evidence that his belief in Q fitted, and perhaps, his noontime Q-belief might be said to be "based on" that old, forgotten evidence. But can a basing relation help Feldman and Conee in this case? Their analysis of well-foundedness explicitly requires that justifying evidential states must be held *at the same time* as the target attitude. But the earlier justifying evidence is no longer possessed, neither consciously nor unconsciously. So evidentialism implies

that Ichabod's noontime belief in Q is not justified (well-founded), an intuitively incorrect verdict. It should be emphasized that in the Ichabod case there is no memory *experience* that triggers his noontime belief in Q, the sort of experience foundationalists typically invoke to account for justified memory-based belief. But many justified memory-based beliefs – especially stored, or non-occurrent, memory-based beliefs -- are unaccompanied by memory experiences. Nonetheless, such beliefs need to be credited with justifiedness. Unless they are justified, they cannot be known -- certainly not on a justificationist view of knowledge. Yet non-occurrent knowledge -- that is, knowledge involving non-occurrent belief -- constitutes the vast bulk of our knowledge at any given moment. If, as theorists, we abandon such knowledge, the skeptic will have won a major (but unearned) victory.

How can evidentialism cope with this problem? It could improve its handling of these cases by abandoning the simultaneity requirement, the requirement that justifying evidence must be possessed at the same time as the belief. But this requirement is a core part of internalism, to which mentalist evidentialism adheres. In any case, how much would it help evidentialism to follow the above advice and abandon the simultaneity requirement? It would still have to add a rule or condition to mark preservative memory as a justification-transmitting feature. It would be extremely *ad hoc* to simply add such a rule or condition without offering a rationale for it. Why, epistemologists are entitled to ask, does this sort of memory process qualify as a justification transmitter? This calls for explanation, and evidentialism has little in its toolbox to offer. What would improve evidentialism's explanatory prospects on this topic is to borrow two basic ingredients of process reliabilism. These ingredients are (1) the epistemological importance of belief-forming or belief-retaining *processes,* and (2) the importance of their *reliability* (or conditional reliability). Preservative memory is a cognitive belief-retaining process that is able to transmit justifiedness from an earlier to a later time. This is why Ichabod is still (substantially) justified in believing Q at noon although he has no evidence at noon that this belief fits. Moreover, reliabilism has a plausible-looking explanation of *why* preservative memory is justification transmitting, namely, its conditional reliability. Belief outputs of preservative memory at later times are mostly true if its inputs at earlier times were true (see Goldman, 1979, for a discussion of conditional reliability).[7] It is not

clear how evidentialism can accommodate any of this within its existing (pure) fittingness story, which is why it needs a "rescue" by reliabilism.

Some readers might not find my treatment of Ichabod persuasive. I stipulated that Ichabod has no other beliefs at noon that provide adequate evidence either for or against Q. But is this plausible, really? A critic might argue as follows. Won't anybody in Ichabod's situation have general beliefs both about the quality of his memory and about the quality of his usual belief-forming processes? Won't he typically believe – and believe justifiably – that his memory is quite good and that he usually forms beliefs in a justifiable fashion? So, if he finds himself remembering that Q and takes account of these background (justified) beliefs, won't this justify him in believing Q *without any reliance on his original evidence*? So there is no need to advert to his forgotten evidence to account for his justifiedness.

The case just described may well be a typical one, but it isn't the only possible case. So let us consider a different scenario. Suppose Ichabod is being treated by a clinical psychologist, who falsely persuades him that his once-robust memory is no longer working well; so he has no reason now for supposing that a stored belief of his (like Q) was accurately transmitted from the past. The epistemological critic imagined in the previous paragraph would then have to say that Ichabod is not justified in believing Q. But is this verdict intuitively correct, given that Ichabod *in fact* reliably retained Q from the past and originally acquired the belief in a perfectly sound fashion? At a minimum, Ichabod's original acquisition and retention of the belief constitute positive forces, or vectors, in support of the justifiedness of his belief. The newly formed beliefs about his powers induced by the psychologist may provide a countervailing epistemic force, a vector that militates in the opposite direction of justifiedness. But that does not negate the fact that his past evidence for Q and subsequent retention of Q are factors *prior* to the current (noon-time) belief that positively affect his current justificational status, something denied by evidentialism.

We might supplement this last point by noticing another way to understand the critic's intuitions without conceding his central claim. His central claim is that the deceived Ichabod is (definitely) *un*justified in believing Q. I am not prepared to concede this (intuitively). On the other hand, I may be prepared to concede a related thesis,

namely, that Ichabod is unjustified in believing *that he is justified* in believing Q. Since Ichabod believes (albeit falsely) that he no longer enjoys the same powers of memory that he once had, he is not justified in believing (upon reflection) that his current belief in Q is justified. In other words, he lacks *second-order* justifiedness with respect to Q. It does not follow from this, however, that he lacks *first-order* justifiedness; i.e., it does not follow that he is not justified in believing Q. As a general matter, being justified in believing that one is justified in believing p does not entail being justified in believing p (JJp $\not\Rightarrow$ Jp). Similarly, being unjustified in believing that one is justified in believing p does not entail that one is unjustified in believing p (~JJp $\not\Rightarrow$ ~Jp). We might concede that (~JJp) is always a *defeater* for Jp, but this is not incompatible with there being defeasible factors working *in favor of* Jp. Indeed, my contention is that Ichabod's sound acquisition of Q in the past and his retention of Q via reliable preservative memory are both factors that work (defeasibly) in favor of his being justified with respect to believing Q. Even if these defeasible factors are overridden by Ichabod's other current beliefs -- a debatable matter -- this would not save evidentialism from the current critique. The provisionally conceded defeat of his justifiedness with respect to Q still allows factors prior to t to be (positively) relevant to the justificational status of his belief at t (i.e., noon), contrary to the thesis of evidentialism.

4. Toward a Two-Component Theory of Inferential Justification

In this section I return to the ecumenical project of proposing that reliabilism should incorporate something like the evidential element emphasized by evidentialism. Earlier versions of process reliabilism tried to make do with the reliability of cognitive processes as the linchpin. To be sure, it has never hesitated to invoke mental *states*, both doxastic and non-doxastic states, in its set of resources (Goldman, 1979, 1986). But mental states have usually been treated as mere inputs and outputs of processes. No fittingness relation between the target belief and these states has figured in the formulations of reliabilism. Now, however, I am inclined to suggest that an incorporation of a fittingness relation into a reformulation of process reliabilism might be a salutary addition. This would generate a *two-component* theory of justifiedness as opposed to a *single-component* theory.

The attraction of a two-component approach seems especially obvious in inferential justification. Let us see how it would work by first replacing the familiar tripartite framework of doxastic attitudes--belief, disbelief, and suspension—with the richer framework involving degrees of credence or subjective probabilities. Next let us assume that for any ordered pair of propositions (P, Q), there is some relation of degree of support or confirmation that takes values on the unit interval. Next, assume that P expresses S's total doxastic evidence relevant to Q, i.e., the total evidence vis-à-vis Q possessed in doxastic form,[8] and that S has no non-doxastic evidence relevant to Q. Finally, assume that the degree of confirmation P confers on Q is N; in other words, C(Q, P) = N. Then it is plausible to hold that doxastic attitude of degree N toward proposition Q is precisely the attitude that *fits* his evidence for Q. In other words, if C(Q, P) = 0.63, then adopting the degree of credence, or subjective probability, 0.63 would *fit* the (total) evidence consisting of P. For an evidentialist about justification, the doxastic attitude a person in the foregoing conditions is *justified* in having vis-à-vis Q is 0.63 (and nothing else).[9]

This account is squarely in the spirit of the fittingness approach; it contains no trace of processism (thus far). The need for a process-based evaluation can be motivated, however, as follows. Shirley is very poor at determining confirmation relations. When she reflects on her total set of beliefs relevant to a hypothesis, she typically draws a blank about the strength of confirmation. She then hazards a wild guess about their degree of support for the hypothesis and forms that degree of belief in it. Proceeding in this fashion, Shirley assigns degree of belief 0.45 to a proposition H. Now consider Madeleine, a highly proficient confirmation theorist. Madeleine has exactly the same evidence vis-à-vis H as Shirley does. She uses her accurate, well-honed skills at determining degrees of support and arrives at the conclusion that her evidence for H is 0.45. She therefore assigns degree of belief 0.45 to H. Now, on one dimension of justifiedness – the fittingness dimension – Shirley and Madeleine's doxastic attitudes vis-à-vis H deserve the same rating. Equally clearly, however, there is another dimension of justifiedness – call it the process dimension – on which their doxastic attitudes merit different ratings. Madeleine's degree of belief is much more aptly, or competently, chosen than Shirley's – despite the fact that they arrive at the same result. On this second

dimension of justifiedness, Shirley's degree of belief is not at all justified or well-founded, whereas Madeleine's degree of belief is very well-founded. A two-factor theory handles this case nicely. But no single-factor theory, of either the purely evidentialist or purely reliabilist sort, can do so. That's a good reason to promote a synthesis of the two.

In their recent paper, "Evidence," Conee and Feldman (2008) anticipate part of what I have just said. They anticipate the objection that assigning a belief or degree of belief as a function of the degree of logical or probabilistic support from one's total evidence does not suffice for justifiedness. They do not, however, explicitly consider endorsing any move toward a process factor to accommodate this objection. Instead they write:

> A person may know some propositions that logically entail some proposition that the person scarcely understands and surely does not know to follow from the things she does know. The logical route from what she knows to this proposition may be complex and go beyond her understanding, or even the understanding of any person. In our view, the person is not then justified in believing the consequence, even though it is entailed by her evidence. It is noteworthy that, to become justified in believing the proposition, she has to learn something new—namely, its logical connection to her evidence. (2008: 94)

Let us set aside the possibility that the subject scarcely understands the target proposition. That feature needs separate attention.[10] The main problem, correctly identified in this passage, is that even entailment by the evidence does not suffice to make the subject justified in believing the target proposition (call it "H"). What is Conee and Feldman's solution? Initially it appears that they seek to solve the problem by adding a requirement that the subject knows the logical connection between the evidence propositions and H. This invites two worries. First, knowledge of a logical (or, more generally, probabilistic) connection does not seem to be necessary for being justified in believing H. They themselves concede this in the following section of their paper. People do not have to learn principles of logic or logical theory (and similarly for probability theory) to come to be justified in believing on the basis of logically or probabilistically supportive evidence.

Second, how would an acquisition of principles of logic or probability help solve the problem? Unless the subject applies appropriate processes, or operations, *to* the newly learned principles to form a belief in H, he will still fail to have a justified belief. At least this is so if we are discussing doxastic justifiedness (well-foundedness), as I am doing here.[11] Where well-foundedness is in question, causal processes of belief-formation (and belief-retention) cannot be brushed under the rug. That is why a two-factor theory, which includes a process dimension, is needed.

5. Ultimate and Derivative Evidence

In this section I consider the kinds of states that qualify as evidential states according to Conee and Feldman and the fittingness relation that is supposed to hold between them and justified doxastic attitudes. Conee and Feldman write:

> The justificatory status of a person's doxastic attitudes strongly supervenes on the person's occurrent and dispositional mental states, events, and conditions. ("Internalism Defended," 2001/2004: 56)

In another statement, which appears in a 2004 "Afterward" to "Evidentialism" (Conee and Feldman, 2004), they write:

> [ES]. The epistemic justification of anyone's doxastic attitude toward any proposition at any time strongly supervenes on the evidence that the person has *at the time* (2004: 101, emphasis added).

Thus, mental states, events and conditions, according to Conee and Feldman, exhaust the evidence (evidential states) that a person ever has, in the preferred sense of evidence they clarify in "Evidence" (Conee and Feldman, 2008). What they call *ultimate* evidence is experiential evidence, including perceptual and memorial experience. Beliefs are also said to be evidence, but not ultimate evidence. "We hold that experiences can be evidence, and beliefs are only derivatively evidence" (2008: 87). Beliefs are "kosher" items of evidence, presumably, because they are mental states and hence internal states. States of knowing, on the other hand, should not be evidential states, because knowing entails truth and truth is not internal. (Conee and Feldman do not subscribe to

Williamson's view that knowing is a mental state.) Strictly speaking, then, it was a little slip when they wrote the passage quoted earlier (from p. 94 of "Evidence") that seems to speak of "knowing" as evidence.

If one puts together the passage quoted above from the 2004 "Afterward" with the statement from the 2008 paper "Evidence" about beliefs being merely derivative evidence, it appears that Conee and Feldman mean to hold that the evidential status of any belief held at time t is derivative from the evidential status of experiential states undergone *at t*. Let us reflect on the implications of this view. Call the total body of evidence – i.e., evidential states – that S possesses at a given time his *concurrent body of evidence* (CBE). What exactly comprises S's CBE at a chosen time t? Presumably, all of his experiential states at t plus all of his beliefs and other doxastic states at t -- or at least the *justified* doxastic states. What I now wish to ask is whether the belief states in CBE are merely redundant. In other words, does the "evidential juice" possessed in CBE reside fundamentally and exclusively in the *ultimate evidence* portion of CBE? Or do some of the CBE beliefs contribute *additional* evidential juice beyond what the ultimate, experiential evidence provides? If we ask what attitude S should adopt toward a given hypothesis H at t, is it necessary to consider S's beliefs in addition to S's experiences? Or is the fittingness or nonfittingness of a possible doxastic attitude toward H determined by the *experiential* portion alone of the CBE? I interpret Conee and Feldman to mean that the experiential portion of CBE suffices to determine whether the fittingness relation holds or does not hold between any hypotheses and the subject's evidence. Call this the *sufficiency of current experience* thesis (SCE thesis). The SCE thesis appears to be what they mean in saying that experience is the *ultimate* evidence. Let us ask, then, whether the SCE thesis is defensible, and whether it helps constitute an adequate theory of justifiedness.

Does the experiential portion of CBE exhaust CBE's evidential power? No, I respond. Many beliefs in CBE can add *extra*, *independently derived* evidential weight to the total evidence beyond the portion of evidence arising from the experiential portion of CBE. Why? Because, in a great many cases, justified beliefs held at time t were originally acquired earlier, as a result of previous experiences. The source of the

evidential power of these beliefs, therefore, may partly reside in those previous experiences, not in CBE experiences. Although the evidential potency of these beliefs may derive from experience, as Conee and Feldman contend, this does not entail that it all derives from *CBE* experience. Thus, the doxastic portion of CBE need not be a merely redundant portion of *CBE*. Its justificational juice may derive from experiences other than CBE experiences.

Conee and Feldman might reply that the mere fact that a CBE belief (justifiedly) originated in past experiences does not imply that it still possesses evidential power at t. It retains such power only if, *at t*, S remembers either the originating experiences or intervening experiences that support the target belief. However, a *memory* at t of earlier experiences is itself an *experience* that occurs at t, and hence belongs to CBE. So the CBE belief acquires whatever evidential power it has from a CBE experience, just as the SCE thesis maintains. In the cases in question, though, the pertinent CBE experience is a memory experience.

This reply is unpersuasive. People frequently fail to recall their original sources of evidence for things they know or justifiably believe. First, they often don't have memory-based *beliefs* about their original sources. More to the present point, they rarely have memory *experiences* of specific perceptual episodes (what psychologists call "episodic" memories) in which they were exposed to relevant observational evidence. It is extremely implausible to claim that for every moment at which a justified belief is stored in your mind, you undergo an episodic memory *experience* of one or more past (perceptual) experiences that constituted your evidence. Five minutes ago I had the stored, dispositional belief that my social security number is such-and-such. Did I also have, at the same time, a memory *experience* of earlier perceptual experiences that provided evidence for this belief? This is dubious. It is especially dubious if "experiences" are restricted to *conscious* mental events. I simply had no such conscious memories five minutes ago. What about *unconscious* memory experiences of the relevant kind? This is what Conee and Feldman would need to invoke to shore up the currently hypothesized line of defense. But what evidence could they adduce for the psychological thesis that whenever a person has a stored factual belief, it is accompanied

by memories of specific perceptual observations of the required sort? Even if such psychological evidence were forthcoming, are such episodic memories *necessary* for the original belief's justifiedness? A psychologically more plausible view is that many beliefs, including my social-security-number belief, are retained via preservative memory from earlier stages of the same belief.[12] An epistemologically more plausible view is that preservative memory of this kind suffices (absent defeaters) to transmit justifiedness from the earlier period. The epistemic role of preservative memory also explains how the evidential status of a belief held at t can partly derive from earlier experiences without positing event memories, at t, *of* such earlier experiences. Thus, the SCE thesis is not well supported.

6. Internalism and the Historicity of Justifiedness

The failure of the SCE thesis and the importance of preservative memory in a satisfactory account of justifiedness pose a serious problem for evidentialism, especially the internalist element in evidentialism. This is so for two reasons. First, the failure of the SCE thesis and the epistemic significance of preservative memory demonstrate the general historicity of justifiedness (see Goldman 2009 for more detail). Although there may be *some* cases in which a doxastic attitude's justificational status is wholly determined by events occurring at its own doxastic decision time, t, *in general* an attitude's justificational status is partly determined by events and states before t. Hence, an attitude's justificational status does not supervene on mental states occurring at t. This is incompatible with Feldman and Conee's evidentialism. According to their principle EJ, it will be recalled, an attitude's justificational status is said to depend on the evidence the subject has *at t*. This is too restrictive. Earlier evidence is also relevant to justifiedness.

It is not just earlier *evidence* that is relevant. Earlier *processes* leading to the belief can also be relevant. Preservative memory is one such process. Suppose S originally formed a justified belief that P based on perceptual experiences of his between t_o and t_1. At a later time, t_n, S also believes that P. Is S's belief that P at t_n also justified? That depends on how S *arrived* at this belief state at t_n. Suppose that the t_n-belief in P is

the product of preservative memory operating from time t_1, and S has encountered no defeating evidence of P during the interim. Then his P-belief at t_n is presumably justified. By contrast, suppose S totally forgets P at some point after t_1, but comes to believe it once again at t_n for entirely frivolous, or spurious, reasons. Then S's P-belief at t_n is unjustified. Thus, which causal process leads to the (dated) belief is crucially relevant.

Evidentialism seems to ignore causal processes. In discussing the evidence that a doxastic attitude must fit in order to be justified, Conee and Feldman (especially in Conee and Feldman, 2008) talk only about beliefs and experiences. Moreover, the evidential status of beliefs is said to derive entirely from that of experiences. Ultimately, then, only experience provides evidence; and it is fittingness with evidence, and only evidence, that confers justifiedness. It appears, then, that processes are left out of the picture. This is unacceptable, as the case of preservative memory illustrates. In principle, an evidentialist might try to evade this problem by declaring preservative memory itself a species of evidence. But this would be a very odd use of the term 'evidence'. Moreover, Conee and Feldman show no signs of wanting to go down this path. Their examples of items (or states) of evidence do not incorporate any such phenomena.

Conee and Feldman might reply that the role of preservative memory is already covered by the basing relation that they admit as relevant in their account of well-foundedness (or doxastic justifiedness). Perhaps, but this is problematic. Clause (ii) of their analysis of well-foundedness speaks of S having doxastic attitude D toward p on the "basis of" evidence E. Does it make sense to view preservative memory as a species of *evidence* on the *basis* of which one might have doxastic attitudes? The things that serve as evidence, under Conee and Feldman's construal, seem to be *states* or *events* with propositional contents. But preservative memory is not a state or an event, although the things it preserves or retains are states or events (or at least the *contents* of such states or events).

There are two other problems facing Feldman and Conee's supervenience thesis (namely, that an attitude's justificational status at t supervenes on the subject's *mental* states at t). Some properties that do not seem to be purely mental also have to figure in

the supervenience base of an attitude's justifiedness. First, fittingness with one's total set of beliefs (at t) is not exactly what determines justifiedness, even waiving the problem of prior acquisition. Only fittingness with *justified* beliefs (and other doxastic states) should determine the justifiedness of a further attitude. Suppose that believing Q is a "proper" inferential response to the total set of beliefs B*. But suppose that no members of B* are justified. All of the beliefs are just products of wishful thinking, or other unsound methods of belief formation. Then a belief in Q (on the basis of B*) does not constitute a justified belief. Justified beliefs cannot be derived from unjustified beliefs. This would not be a problem for the mentalist supervenience thesis if the justifiedness property of a belief held at t supervened on mental states obtaining at t. But we have already seen that this is not so. The justificational status of an attitude held at t supervenes on what happens both at t and prior to t.

A second problem is the fittingness relation itself. Fittingness is not a mental state of affairs. According to evidentialism, given the composition of a particular CBE at time t, believing H either fits or does not fit the CBE. Whether the fittingness relation is instantiated by the ordered pair (B(H), CBE) does not depend on whether the subject (or anybody else) *thinks* it is instantiated. Thus, the justificational status of the subject's believing H depends on an *instantiation fact* (as it is called in Goldman, 2009), which is not a mentalistic fact. Here is yet another problem for the mentalistic supervenience thesis about justification.

Another lurking problem is one for internalism more generally. Fittingness is a seemingly attractive relation for internalism because it seems to be an internal relation, hence a welcome tool or resource for this type of approach. Consider what the internalist Chisholm wrote about the nature of justification: "the concept of epistemic justification is … *internal* and *immediate* in that one can find out directly, by *reflection*, what one is justified in believing at any time" (1989: 7). The fittingness relation looks like it would find a comfortable home in Chisholm's internalist conception. Surely, one can tell by reflection alone (ostensibly an a priori process) whether a given attitude toward hypothesis H does or does not fit with one's current CBE. After all, can't one tell by

reflection what one's current CBE *is*, and can't one tell by reflection whether the selected attitude toward H does or does not fit one's CBE?

Unfortunately, this simple picture is incorrect. One cannot tell by *reflection* (if reflection is supposed to be an a priori process) what one's current mental states are. One can only determine the composition of one's total mental state by *introspection*, which is not an a priori process (see Goldman, 1999b). (However, perhaps Chisholm meant to include introspection under reflection.) Second, and more importantly, one's CBE – one's total body of *evidential* states -- is not identical with one's total body of *mental* states. The total body of evidential states includes only one's *justified* beliefs, not the unjustified ones. Although one can tell by introspection which beliefs one has (at least conscious beliefs -- stored beliefs presumably are not introspectible), one cannot tell by introspection which ones are *justified*. The justificational status of one's beliefs, as we have seen, partly depends on their history of acquisition and confirmation (prior to t). None of this is accessible to introspection at time t. If one cannot tell by introspection (at t) which of one's beliefs are *evidential* states, then one cannot tell by reflection (at t) whether the fittingness relation holds between one's CBE and the selected attitude in question. So, instantiations of the fittingness relation are not generally detectable (at the time of attitude selection) by reflection alone.

It emerges from these points that the fittingness relation should hold less appeal to internalists than it initially appeared. This may not matter to Feldman and Conee, who in do not accept Chisholm's accessibilism. For accessibility internalists, however, the above discussion is quite pertinent, insofar as it suggests that the fittingness relation is a less promising piece of armor for defending their cause than they might imagine.

This point is independent of the many complications that would be encountered in trying to work out a determinate fittingness relation that meets Feldman and Conee's desiderata – an enormous problem given the special difficulty of balancing the respective roles for current experience and current beliefs. However, since I am tentatively prepared to sign on to the notion of a fittingness relation (as indicated in section 4), I won't say that it cannot be done. The question I have mainly addressed in this section is whether a

well-specified fittingness relation would be such a major victory for internalism. I have argued in the negative.

7. Experiential Evidence: Toward a Two-Component Theory

As indicated at the outset, one aim of this paper is to explore the prospects for a hybrid theory of justification that combines the strengths of reliabilism and evidentialism. Thus far, however, rather limited steps have been taken toward embracing evidentialist themes. The main such step (in section 4) was to advocate a two-factor theory of inferential justification, in which one factor is a belief's fittingness relation to the subject's *doxastic* evidence. Should reliabilism go a step further and acknowledge an analogous role for fittingness with *non-doxastic* evidence, for cases of non-inferential justification? This is what Comesana recommends in his proposed synthesis. He motivates this move in response to the oft-cited counterexamples to reliabilism offered by BonJour (1985) and Lehrer (1990). The examples of Norman the clairvoyant (BonJour) and Mr. Truetemp (Lehrer) purport to demonstrate the non-sufficiency of reliability for justifiedness. Comesana's diagnosis of these cases is that the subject lacks any *evidence* – especially *experiential* evidence – on which his belief is based. Nonetheless, the beliefs in these cases are all reliably caused. So reliabilism must classify them as justified -- a mistaken classification according to most commentators. Comesana therefore proposes a strengthened sufficiency condition according to which a belief is justified if it is caused by a reliable process *that includes some evidence*. We earlier acknowledged a need for doxastic evidence in the case of inferential justification. Why not join Comesana in taking the further step of requiring non-doxastic evidence, i.e., experiential evidence, for non-inferential justifiedness? This would handle the clairvoyance and Truetemp cases and would add more heft to the proposed synthesis of reliabilism and evidentialism.

An intriguing counter-argument against this concession to evidentialism is made by Jack Lyons (2009). Lyons argues that experiential grounds (or evidence) play no essential role in the structure of justification. In particular, he holds that there can be -- and there are -- immediately justified ("basic") perceptual beliefs that do not get their justification from perceptual *experiences* that serve as grounds. Instead, these perceptual

beliefs get their justifiedness from their being (A) basic and (B) the products of reliable processes. He defends this position by appeal to both possible and actual cases of justified perceptual belief without consciously experienced grounds. One hypothetical case he adduces is zombies,[13] who have justified perceptual beliefs in the absence of any conscious perceptual experience. Another hypothetical case is Block's (1995) super-blindsighter, who has visual beliefs without visual experiences. An actual case is one of J. J. Gibson's (1966) examples of "sensationless perception." This example involves blind people who possess an obstacle sense: they detect obstacles like walls and chairs without having any associated conscious sensations. They tend to think they are picking up information somehow through the skin of the face, when in truth the information is coming through the ears as a subtle form of echolocation (Lyons, 2009: 52). Lyons regards the perceptual beliefs associated with this obstacle sense as justified.[14]

Lyons also offers an account of basic beliefs that requires them to be outputs of a certain type of modular system -- what he calls a "primal system" -- but makes no requirement concerning an experiential ground. His way of handling the Norman and Truetemp cases is quite different from Comesana's. He says that their beliefs are not basic because basic beliefs must be the products of a primal system and no such system would underpin Norman's clairvoyance or Truetemp's temperature detection. Under Lyons's theory, reliability suffices for the justifiedness of basic beliefs; but Norman's and Truetemp's beliefs do not qualify as basic. I shall not try to decide if Lyons's alternative strategy ultimately succeeds.

However, even if we agree with Lyons that experience is inessential for basic perceptual justifiedness and that a different explanation of the clairvoyance and Truetemp cases is possible (without abandoning reliabilism), this does not demonstrate that all is well with reliabilism's treatment of experience. Few epistemologists would deny that people have both perceptual and memorial experiences, and it is reasonable to expect such experiences to play a distinctive role in justification. The one traditional epistemology that ignores or downplays the evidential role of experience is coherentism, and its standing in epistemology takes a hit precisely because of this feature. Reliabilism shows no comparable opposition to a role for experience, but it does little to highlight or

acknowledge such a role. Its silence about experiential evidence is, at a minimum, a noticeable lacuna.

Here is a specific argument to motivate a substantive role for experiential states in a theory of justification, an argument from defeaters. Sidney inferred it would be sunny this afternoon from what he read in this morning's newspaper. On the basis of this evidence, he continues to believe an updated version of this proposition – namely, that it is sunny right now (in the middle of the afternoon) -- despite the fact that he is walking in the middle of a rainstorm. Surely his current perceptual experience is a defeater for this belief; he is not justified in believing that it is sunny right now. But he has used no belief-forming or belief-revising process that takes this perceptual experience as an input; he just ignores this experience. It appears, then, that there is no *process* we can appeal to – at least no *instantiated* process – to account for the defeat of his sunniness belief. If we want to say what defeats Sidney's current justification for his sunniness belief, the obvious candidate is his perceptual experience. Moreover, it is natural to say that this justification-undermining experience is a piece of evidence he possesses.[15]

In section 4 I offered a two-factor approach to inferential justification in which doxastic evidence occupied a central role. A similar move is in order with respect to experiential evidence, in light of the foregoing argument from defeat (among other things). So, an analogous two-component approach to non-doxastic justification (that is, justification that rests on non-doxastic states) is in order, in which perceptual and memorial experience would occupy a central role. Let us concentrate on perceptual experience.

To construct a two-factor theory of perceptual justification we need an account of perceptual "fittingness" to provide the first factor. A preliminary question is how to construe perceptual states. Presumably they are not *doxastic* states; but might they still be *contentful*, or *representational,* states? To remain in the vicinity of Feldman's (2003) approach to perceptual justification, we would have to interpret them as being representational, or contentful, states. Otherwise, there could be no *closeness of content* between a perceptual state and beliefs that it justifies (or helps to justify). If perceptual

states have content, of course, this might scare up the ghost of the Sellarsian dilemma. For justification to derive from a contentful state, doesn't that state itself have to be justified? And if it must be justified, how does it attain this status?

This initial problem is not so daunting. Both philosophers of mind and cognitive scientists commonly treat perceptual states as representational states. This does not mean that they are doxastic (belief-like) states, and if they are not doxastic states they do not need to be justified themselves in order to confer justifiedness on other states. But how can they confer justifiedness if they don't possess it themselves? As evidentialism says, they can confer justifiedness without being justified because they are evidence. In virtue of what are they evidence? As proposed in section 1, we can provide an answer congenial to reliabilism: they are evidence (or evidential states) because they are (commonly) *reliable indicators* of the subject's environment. Perceptual states are produced by systematic relations between properties of the environment, properties of a perceptual medium (light waves, sound waves, etc.), and perceptual systems that pick up information from these sources and transmit it to the brain. In favorable circumstances, the resulting conscious states are systematically correlated with tokenings of relevant properties of (objects in) the environment, and hence are evidence for such tokenings.

Granted that these states are evidence (hence capable of conferring justification), how can beliefs about the environment *fit* this evidence? A feature of the ecumenical account of inferential justification proposed in section 4 was acceptance of fittingess as one component of the story. Is the fittingness relation applicable here as well, in the perceptual case? As remarked above, if the approach considered here is to mesh with the one sketched in Feldman 2003, it would have to involve some sort of content matching, content overlap, or content "implication" between perceptual states and beliefs (or other doxastic states). How good are the prospects for such a contentful relationship? It is widely held that perceptual states have a different species of content than doxastic states do: nonconceptual rather than conceptual content. Is this correct, and would it preclude serious prospects for a suitable contentful relationship between the two kinds of states?

There are many different interpretations of the distinction between conceptual and nonconceptual content. It is not clear, however, that any of them definitely implies that these are two distinct *kinds* of content. Jeff Speaks (2005) offers a general critique of this conclusion. Here is how he argues, for example, against just one defense of this idea. Many writers appeal to the richness, or fine-grained character, of experience to show that the contents of experience are nonconceptual. A thousand words, it may be said, would hardly begin to convey a complete description of how the world appears to you in a brief perceptual state. Speaks responds that even if experience is far more detailed and full of information than could be captured in a single thought, or even a lifetime of thoughts, this does not show that the information given in perception is of a different *kind* than the information represented by a belief. It only shows, at most, that there is more of it.

Without resolving this issue, let us turn to the question of fittingness. How could a doxastic attitude toward a proposition "fit" an experiential state, assuming that both have content (whether the same or different kinds of content)? In the case of inferential fittingness, we proposed that the relevant relation is that the truth of the contents of the premises should *make probable* the truth of the content of the conclusion. Perhaps the same idea can work here. Perhaps we can say that a belief in proposition p fits the evidence provided by experiential state E just in case the truth of the content of E *makes (highly) probable* the truth of p.

We can illustrate this idea with the help of a particular proposal in vision science for how beliefs about visible objects are formed based on visual representations of shape.[16] According to this proposal, the visual system represents object *shapes* in terms of parts constructed out of geometrical "ions" (elements), called *geons*. Geons are members of a family of representations of volumetric shapes that can be modeled as generalized cones, i.e., volumes swept out by a cross-section moving along an axis. A typical geon is a cylinder. Groups of geons can be combined by relating several of them to one another in ways familiar from tinkertoy sets. A combination of geons plus selected relations among them (e.g., "side-attached") can be recognized as familiar kinds of objects, such as a chair, a giraffe, a mushroom, or a pail.

Assume now that geonic combinations are a type of (high-level) visual representation, and that, in normal circumstances, a specific geonic combination G is tokened only when a giraffe is present. Then the truth of a G-token's content makes it highly probable that the proposition "A giraffe is in the vicinity" is true. Hence, it is fitting for a subject to believe the latter proposition when a G-representation is tokened. This, then, is an adaptation of the evidentialist notion of fittingness to the case of visual experiential evidence, thus making fittingness one suitable factor, or component, of experience-based justification.

If we stop here, of course, we have a purely evidentialist account of experiential justification. It doesn't resemble an evidentialist-reliabilist hybrid of the sort I have been touting. Must we add a process-reliability component to obtain a more satisfactory account? Yes. Otherwise, we won't cover all the necessary bases. In particular, we won't have a satisfactory account of doxastic as opposed to propositional justifiedness (i.e., well-foundedness) in this domain.

Obviously, to have a justified belief that object x is an F, where the justification arises from shape perception, it is not enough that a geonic F-representation is tokened in the subject's head while he forms the indicated belief. This condition might be met even by the novice bird-watcher in Feldman's example (see section 2 above) who should not be credited with a justified belief. The novice bird-watcher might even undergo a tokening of the same geonic representation of the pink-spotted fly-catcher as the expert bird-watcher undergoes. In that case, the fittingness condition will be satisfied by both of them if it is satisfied by one. But, intuitively, it is still possible that one should be justified and the other unjustified (as Feldman recognizes).[17]

Biederman (1987, 1990), the leading proponent of the geonic approach to visual object-recognition, gives us the materials for explaining the difference between the two bird-watchers in terms of the geon theory. Biederman says that subjects construct "object-models" for a large number of object names in English, where each "object-model" is a particular geonic configuration of specific geons and relations between them. Many object types require several models each, both because different models are

required for different viewing orientations and because some names or concepts have more than one configuration associated with them. For example, lamps come in different configurations, and there are several types of pianos (grand, baby-grand, upright, and spinet) each of which has a distinctive geonic configuration. Biederman hypothesizes that people store (geonic) object-models in memory where they are linked to object names or concepts. Then, when they have a visual experience, they automatically try to *match* a geonic configurations tokened in visual consciousness with one of the geonic configurations stored in memory. When such a "match" is secured, the system *recognizes* (or judges) the perceived object as an F, where F is the label or concept paired with the recovered object-model.[18]

Presumably, this matching process would be the kind of process used by the expert bird-watcher, who has a stored geonic object-model for the pink-spotted fly-catcher. If his classification of the bird as a pink-spotted fly-catcher is the product of a matching process of the indicated kind -- more specifically, a matching process with a high threshold or standard for a match (see note 18) -- he uses a very reliable process of classification, or belief-formation. So this would satisfy a process-reliabilist condition of the kind I mean to incorporate into our hybrid account. The novice bird-watcher, by contrast, would not be using such a reliable process. His lack of training suggests that he has not constructed a suitable object-model for the pink-spotted fly-catcher, and therefore does not secure much of a "match" at all between his visual experience and a stored object-model of the fly-catcher. Hence, he is just guessing when he classifies the bird as a fly-catcher, and guessing is not a reliable process.

I propose, then, in parallel with the proposal of section 4, that experiential justifiedness is a function of two factors, or components. One factor says that belief in a proposition is *prima facie* justified in experiential terms only if the belief fits with the subject's current experiential evidence.[19] The second factor says that an experience-based belief is justified only if it is the product of a reliable experience-based process. The two-factor theory says that an experiential belief is fully justified (doxastic considerations aside) only if it is justified in terms of both factors. This two-factor approach to experience-based justifiedness, which marries the requirement of evidential

fittingness with production by a reliable process, has significant attractions, I have argued.

8. The Source of Epistemic Principles (Best Explanationism)

In their recent paper "Evidence," Conee and Feldman (2008) rightly criticize Chisholm's approach to epistemic principles. Chisholm denies that epistemic principles derive from any more fundamental source, and he does not seek a unified account of them. By contrast, Conee and Feldman contend that if perceptual and memorial experiences are justifying, there must be something about them that makes this the case. And if religious experience is similarly justifying, this must be so for the same underlying reason. Here we are in complete agreement. Conee and Feldman are pointing to what I call (Goldman, 2009) a *criterion* or *rationale* of justificational rightness. An adequate theory of justifiedness should specify the content of this criterion or rationale. I endorse the externalist idea that what rationalizes various epistemic rules or principles has to do with the proficiency or efficacy of the principles at meeting such goals as acquiring true belief and avoiding error. A generalization of this idea for degrees of belief is offered in Goldman (1999) under the label "veritistic value" (although that theory is not directed at justifiedness per se). Conee and Feldman, of course, seek an internalist criterion or rationale. The one they propose in "Evidence" is the criterion of "best explanations." Here is how they put it:

> We believe that the fundamental epistemic principles are principles of best explanation. Perceptual experiences can contribute toward the justification of propositions about the world when the propositions are part of the best explanation of those experiences that is available to the person. Similarly, the truth of the contents of a memory experience may be part of the best explanation of the experience itself. Thus, the general idea is that a person has a set of experiences, including perceptual experiences, memorial experiences, and so on. What is justified for the person includes propositions that are part of the best explanation of those experiences available to the person. Likewise, one's inferences justify by identifying to one further propositions that either require

inclusion in one's best explanation for it to retain its quality or enhance the explanation to some extent by their inclusion. (2008: 98)

Unfortunately, there are a number of counterexamples to this approach, plus the worry that best explanation is itself subservient to a different rationale, namely true belief acquisition. I turn first to the counterexamples. It should be emphasized that the success of even a *single* counterexample among the three offered below would suffice to make my critical case. *All* the counterexamples would have to be successfully resisted to overcome this line of criticism.

(A) *Introspection*. How would Conee and Feldman subsume a principle of justifiedness for first-person-current-mental-state beliefs under the explanatory inference model? The most promising tack is to posit an introspective principle of evidence analogous to a visual principle of evidence. In the latter case, a principle might say: "Having a visual experience V(E) is *prima facie* evidence for E's being the case." This can subsumed under the best explanation mode by indicating that having, say, a visual experience as of a mouse running across the floor is best explained by its being the case that a mouse *is* running across the floor. Similarly, it might be proposed, what gives me *prima facie* evidence that I now undergoing a tickle sensation is that I am having an introspective experience as of a tickle sensation. This would be subsumed under the best explanation theme by indicating that having an introspective experience as of a tickle sensation is best explained by its being the case that I *do* have a tickle sensation.

The viability of this approach depends, however, on there being introspective experiences of mental events or states of affairs distinct from the mental states that are allegedly introspected. Is there an introspective experience of a tickle sensation *in addition to* and *distinct from* the tickle sensation itself? This is extremely doubtful. I don't doubt that there are tickle experiences. And I don't doubt that there are (self-reflective) beliefs that one is having a tickle experience. What I doubt is that there is any "intermediate" mental event that stands between the "objective" state of affairs in question (here, the tickle sensation) and the belief in this state of affairs in the same way that a visual experience of a mouse is an "intermediate" event that stands between the

mouse scurrying across the floor and the belief that a mouse is scurrying across the floor. If there is no such intermediate introspective experience, then what is the explanandum (item to be explained) that the occurrence of the tickle sensation (the explanans) supposedly explains? I deny that there is anything that can function in the guise of a suitable explanandum, thereby casting into doubt a best-explanationist model of introspective justification.

(B) *Preservative Memory.* As reviewed above, a satisfactory theory of justifiedness needs a principle of preservative memory. If S was justified in believing P earlier, and S retains her belief in P now via preservative memory, then S is prima facie justified in believing P now. Like introspection, however, preservative memory lacks any type of mental experience or episode that invites explanation. In particular, there is no (conscious) act of "recollection" that invites explanation. An epistemic principle that covers preservative memory, then, cannot be rationalized by the best-explanation approach.

(C) *Arithmetic Inference.* I think there are two squirrels on my deck, and I think there are two birds. So I infer that there are (at least) four animals. Presumably, this arithmetic inference is justified. Is it a case of explanatory inference? Surely not. How does there being four animals *explain* there being two squirrels and two birds? It doesn't. Still, here is a justified belief that some epistemic principle must cover. But that principle, in turn, cannot be grounded in terms of best explanation.

In each of these cases, the best-explanation account does not work. (Or, more cautiously, it does not work in some of them.) At the same time, each is very plausibly accommodated by process reliabilism. The belief that one has a tickle sensation is justified because introspection is a reliable process. Preservative memory contributes to justification because preservative memory is a conditionally reliable process.[20] Arithmetic inference (of the right sort) is justifying because it is a conditionally reliable process. Thus, process reliabilism does a better job of accounting for the justifiedness "data" than best-explanationism.

In addition to these counterexamples, we should reflect on why inference-to-the-best-explanation is often a good pattern of non-deductive inference. I submit that it is a good pattern (where it is) because it is conducive to true belief. So, even if the inference-to-the-best-explanation rationale were accepted, there is still a deeper rationale of an externalist variety. In defending this thesis, I don't wish to appeal to the argument that no explanation genuinely *explains* anything unless it is true. That is doubtless correct in one sense of 'explain'. However, people like Conee and Feldman who invoke inference-to-the-best-explanation must mean 'explain' to be understood in some non-truth-entailing sense of 'explain' (whatever that sense is, exactly). As dyed-in-the-wool internalists, they would not want a fundamental rationale of theirs to appeal directly to truth considerations. Suppose, then, that "superior explanatoriness" is not *defined* in terms of truth-conduciveness. Nonetheless, I submit, superior explanatoriness strikes us as a rationalizing property of a type of inference only because superior explanatoriness is, in general, an excellent *indicator* of truth. In other words, even if we accept superior explanatoriness as a mark (only one mark, not a universal mark) of justification, its being such a mark derives from its correlation with truth-conduciveness. The most fundamental principle of epistemic justification, then, is truth conduciveness -- as reliabilism, of course, maintains.*

References

Biederman, I. (1987). "Recognition-by-Components: A Theory of Human Image Understanding," Psychological Review 94: 115-147.

Biederman, I. (1990). "Higher-Level Vision," in D. Osherson, S. M. Kosslyn, and J. M. Hollerbach, eds., An Invitation to Cognitive Science, vol. 2, Visual Cognition and Action (pp. 41-72). Cambridge: MIT Press.

BonJour, L. (1985). The Structure of Empirical Knowledge. Cambridge, MA: Harvard University Press.

Chisholm, R. M. (1942). "The Problem of the Speckled Hen," Mind 51: 368-373.

Chisholm, R. M. (1989). Theory of Knowledge, 3rd edition. Prentice-Hall.

Comesana, J. (2002). "The Diagonal and the Demon," Philosophical Studies 110: 249-266.

Comesana, J. (forthcoming). "Evidentialist Reliabilism," Nous.

Conee, E. (1988). "The Basic Nature of Epistemic Justification," The Monist, Reprinted in Conee and Feldman, Evidentialism (2004).

Conee, E. and Feldman, R. (2001/2004). "Internalism Defended," in H. Kornblith, ed., Internalism and Externalism, pp. 231-260. Malden, MA: Blackwell. Reprinted in Conee and Feldman, Evidentialism (2004).

Conee, E. and Feldman, R. (2004). Evidentialism: Essays in Epistemology. Oxford: Oxford University Press.

Conee, E. and Feldman, R. (2008). "Evidence," in Q. Smith, ed., Epistemology: New Essays. OUP.

Feldman, R. (1985). "Reliability and Justification," The Monist 68: 159-174.

Feldman, R. (2003). Epistemology. Upper Saddle River, NJ: Prentice-Hall.

Feldman, R. and Conee, E. (1985/2004). "Evidentialism," Philosophical Studies 48: 15-34. Reprinted with an afterward in Conee and Feldman, Evidentialism (2004).

Gibson, J. J. (1966). The Senses Considered as Perceptual Systems. Boston: Houghton Mifflin.

Goldman, A. I. (1979). "What Is Justified Belief?" in G. Pappas, ed., Justification and Knowledge. Dordrecht: Reidel.

Goldman, A. I. (1986). Epistemology and Cognition. Cambridge, MA: Harvard University Press.

Goldman, A. I. (1999a). Knowledge in a Social Wortd. Oxford: Oxford University Press.

Goldman, A. I. (1999b). "A Priori Warrant and Naturalistic Epistemology," in J. Tomberlin, ed., Philosophical Perspectives, vol. 13. Cambridge, MA: Blackwell.

Goldman, A. I. (2008a). "Reliabilism," in E. N. Zalta, ed., Stanford Encyclopedia of Philosophy (Fall 2008 Edition), URL = http://plato.stanford.edu/archives/fall2008/entries/reliabilism/.

Goldman, A. I. (2008b). "Immediate Justification and Process Reliabilism," in Q. Smith, ed., Epistemology: New Essays, pp. 63-82. Oxford: Oxford University Press.

Goldman, A. I. (2009). "Internalism, Externalism, and the Architecture of Justification." The Journal of Philosophy.

Kelly, T. (2006). "Evidence," in E. N. Zalta, ed., Stanford Encyclopedia of Philosophy (Fall 2008 Edition), URL = http://plato.stanford.edu/archives/fall2008/entries/evidence/.

Kim, J. (1988). "What Is Naturalized Epistemology?" in Tomberlin, J. Philosophical Perspectives 2. Atascadero, CA: Ridgeview.

Lehrer, K. (1990). Theory of Knowledge. Boulder, CO: Westview Press.

Lyons, J. (2009). Perception and Basic Beliefs: Zombies, Modules, and the Problem of the External World. New York: Oxford University Press.

Marr, David (1980). Vision: A Computational Investigation into the Human Representation and Processing of Visual Information. San Francisco: W. H. Freeman.

Speaks, J. (2005). "Is There a Problem about Nonconceptual Content?" Philosophical Review 114: 359-398.

Stalnaker, R. (1978). "Assertion," in P. Cole, ed., Syntax and Semantics, vol. 9, Pragmatics. New York: Academic Press.

Williamson, T. (2000). Knowledge and Its Limits. Oxford: Oxford University Press.

[1] See Goldman (1979, 1986), Feldman and Conee (1985), Conee and Feldman (1998), Goldman (1999), Conee and Feldman (2001).

[2] A commitment to defining 'knowledge' in terms of 'justification' is clear at least for Feldman, who writes: "knowledge requires justified true belief that does not essentially depend upon a falsehood" (Feldman, 2003: 37). I know of no reason to suspect that Conee disagrees, at least about justification being definitionally necessary for knowledge.

[3] That Conee and Feldman equate well-foundedness with doxastic justifiedness was indicated in personal communication with Conee and affirmed by Feldman at a 2008 conference at Rochester where Feldman served as commentator on an early version of this material.

[4] By "non-doxastic fittingness" I mean fittingness of a belief to one or more non-doxastic mental states, for example, perceptual experiences. This has nothing to do with the distinction between propositional and doxastic justifiedness.

[5] I don't mean to dismiss content matching or fittingness entirely in an account of non-inferential justifiedness. I return to this topic in section 7 below.

[6] It may be contended that Ichabod's noontime belief in Q is not *as justified* as his belief in Q was when it was originally formed and he was in possession of the original evidence. This point may be conceded. Nonetheless, given the history described, the noontime belief has a substantial measure of justifiedness, and that measure of justifiedness cannot be explained by evidentialism, as far as I can see.

[7] I am not assuming that preservative memory is *defined* as memory with mostly accurate retention of believed content. Rather, I assume it's a contingent fact that what is held in memory from previous periods is generally "true to" what was there earlier. (Of course, this does not preclude a large amount of forgetting.)

[8] In saying that he possesses evidence in "doxastic form," I mean that he possesses evidence in virtue of either believing categorically or having weaker degrees of credence in propositions that are confirmationally relevant to Q. Plausibly, these propositions don't constitute evidence unless his doxastic attitudes toward them are justified (at least semi-justified). I won't try to settle the latter issue here.

[9] Actually, there are more problems lurking here. What dictates the appropriate degree of credence in Q is not only the strength of the confirmational relation between the evidence *propositions* and Q, but also.how strong are the (justified) credal attitudes in these various evidential propositions. (I am assuming that they don't all have to be 1.0, or flat-out belief. I do not feel that it is *my* job to solve this problem. It is a problem that squarely confronts evidentialism, so the main duty of solving it falls on the shoulders of evidentialists. My proposal in this section is that, assuming this problem can be solved, we will *then* have on our hands one kind of justificational factor that could and should be added to the reliable-process factor that reliabilism traditionally emphasizes.)

[10] In any case, it isn't a problem for doxastic justifiedness, the variety of justifiedness on which I am concentrating for now. Unless the subject understands the proposition, he does not believe it, so the issue of doxastic justifiedness does not arise.

[11] It is also relevant, however, to *propositional* justifiedness, as the selection problem discussed in section 2 reveals. The reason that Maurice is not *propositionally* justified in believing that the table is 12 years old is that he lacks any reliable cognitive process that would take his visual appearance as input and yield a belief that the table is 12 years old as output.

[12] This is the kind of "factual" memory that psychologists refer to under the label "semantic" memory (in contrast to "episodic" memory).

[13] Lyons defines "zombies" in a weaker sense than usual in philosophy, as beings that are as psychologically similar to one of us as possible consistent with their lacking conscious experiences (2009: 51-52).

[14] Lyons provides a number of other examples drawn from perceptual psychology. The reader should consult his book for detailed exposition and interpretation of these examples (2009: 52-59). In addition to perceptual cases where there is room for doubt that justifiedness depends on experiential evidence or grounds, I would adduce the case of introspective beliefs. I can be justified in believing that I am now in this or that mental state, but there is no additional mental state (distinct from the target state itself) that serves as ground or evidence for such a belief.

[15] In my earlier formulations of (pure) reliabilism (Goldman, 1979, 1986), I tried to accommodate the problem of defeat or undermining by invoking reliable processes that *could* or *should* have been applied to the undermining mental states, processes which, if used, would have prevented the relevant belief from being formed. But these attempts to accommodate defeat in terms of the process framework were rather strained, I would now say. They were not straightforwardly inadequate, but the treatment can definitely be improved upon by incorporating evidential states explicitly into the theory, as I now propose. .

[16] Not all vision scientists would sign on to this approach, but it has received considerable support and development, especially by Irving Biederman (1987, 1990), and it builds on the very popular ideas of David Marr (1980). However, it is used here mainly because of its vividness, not because of its empirical adequacy as presently judged by the research community.

[17] The bird-watching example makes for a poor fit with the geon theory of object recognition, for the obvious reason that colors of body parts constitute a major feature of bird identification, and colors are not included in the shape-oriented geon theory. For present purposes, however, let us pretend that all of the relevant features involved in bird identification are shape features.

[18] As used here, the term "match" does not refer to a perfect correspondence between all parts of the visually experienced geonic configuration and the complete object model stored in memory. Perfect correspondence may be atypical in perceptual experience because not all parts are usually in view, either because of the object's orientation to the viewer or because of occluding objects. "Match" here refers to a correspondence with respect to *enough* elements of the two geonic configurations to meet the threshold, or criterion, of matching that the subject's psychological system employs on the occasion in question. This threshold can vary from occasion to occasion. Of course, threshold variation can affect the reliability of the matching process upwards or downwards. But Biederman argues, on theoretical grounds, that even a fraction of the elements of a full geonic configuration often suffices to obtain a reliable identification of an object.

[19] I say "prima facie" justified because one might have non-experiential evidence from background beliefs that also affects one's *ultima facie* evidence for believing the proposition. The problem of how to balance conflicts between experiential and non-experiential evidence is not on the present agenda.

[20] As defined in Goldman (1979), a belief-dependent process is conditionally reliable just in case it yields a high ratio of output truths in cases when all of its belief inputs are true.

* I have received valuable comments from Holly Smith, from Richard Feldman and other discussants at the 2008 Rochester graduate epistemology conference, and from audiences at other venues where earlier versions of this paper were presented, i.e., Brown University, Union College, and the Metropolitan University of Mexico City.