

Immediate Justification and Process Reliabilism

Alvin Goldman
Rutgers University
[Final version: October 13, 2006]

1. Introduction

A central issue in contemporary epistemology is whether there is a species of (prima facie) justification that is immediate, direct, basic, or foundational. It is puzzling whether and how immediate justification could arise. This is perhaps the core issue that divides foundationalists from coherentists. An older conception of immediate justification is that some beliefs are capable of *self*-justification; they don't need anything *else* to justify them.¹ This conception of immediate justification is no longer very prevalent. Current defenders of immediate justification usually assume that a directly justified belief is rendered justified by some state of affairs distinct from it. What is special to the direct mode of justification conferral is that it doesn't involve *other* justified beliefs (or states or propositions) as justifiers. Deniers of immediate justification doubt that this occurs, or even makes sense.

In general, questions about justification can be given either a *doxastic* or a *propositional* formulation. A doxastic formulation presupposes that the epistemic subject has a belief of interest, and the question is whether this belief is justified. A propositional formulation does not presuppose that the subject believes the target proposition. It asks whether she is in a "position" such that it *would* be epistemically appropriate for her to believe it, whether or not she does so. We shall address both types of formulation.

The issue of immediate justification can be resolved into three questions:

(Q1) Is there a species of justification that is immediate, direct, or basic? In other words, are some beliefs or propositions made (prima facie) justified in virtue of states of affairs, processes, etc. that confer justification without themselves being justified?

If question Q1 is answered in the affirmative, the next question is:

(Q2) How is immediate justification conferral possible? There are well-known arguments, initially plausible-looking arguments, against the very possibility of unjustified justifiers. What is wrong with those arguments?

Assuming that challenges to the very possibility of unjustified justifiers can be met, the final question is:

(Q3) What types of states or process can serve as immediate justifiers, and what relations do they bear to appropriate beliefs or propositions by virtue of which justification is conferred on those beliefs or propositions and not others?

Here is a sample dilemma that appears to challenge the possibility of basic justification. Initially it seems that anything that makes a proposition or belief justified must itself be justified. This pattern is instantiated in the paradigm case of inferential justification. An inferred proposition is justified (for the subject) only if the premises or beliefs from which it is inferred are themselves justified (for that subject). But if the justifiedness of a proposition or belief depends on the justifiedness of another belief (or belief-like state), then the justifiedness of the target belief or proposition isn't basic or immediate after all. Immediate justification can arise only from states, facts, conditions, processes, or the like that aren't themselves justified, i.e., from unjustified justifiers. At a minimum, immediate justification arises only from states, facts, and so forth without *depending* on their justificational status (though they might *have* such status). But how can unjustified states or conditions confer justifiedness on a proposition or belief?

Another challenge to the possibility of immediate justification is prompted by a second feature of inferential justification. In inferential justification, the source of justification is prior *contentful* (i.e., representational) states of the subject. The contents of these states are what fix the further propositions that are inferentially justifiable for the subject. What shall we say with respect to content for putative cases of direct, non-inferential justification? Must the unjustified justifiers themselves be contentful, representational states? If so, are the only propositions on which they confer justification the contents of the conferring states? If unjustified justifiers can be contentless, non-representational states, how do specific propositions or beliefs get selected as ones made justified by these contentless justification-conferrers?

There seem to be three possible answers to this conundrum:

- (A1) Immediate justifiers are contentful states of affairs (presumably, states of the epistemic agent, S).
- (A2) Immediate justifiers are non-contentful states of affairs.
- (A3) Immediate justifiers can be either contentful states or non-contentful states (or perhaps a combination of the two).

In the last few years several epistemologists have focused their attention on this nest of problems. I shall consider treatments of the topic by four epistemologists: Richard Feldman (2003), James Pryor (2005), Michael Huemer (2001), and Peter Markie (2005). Each offers helpful explorations of the territory, but none of them, in my view, offers a satisfactory positive solution. Markie, for example, treats the problem as an unsolved "mystery". I shall suggest that process reliabilism offers the most promising solution to the problem, or nest of problems. Feldman and Markie will be unhappy with this proposal, because they are on record as finding serious or fatal problems with reliabilism. I won't attempt, however, to address any of the familiar problems for reliabilism, such as the generality problem.² This article is limited to a defense of the thesis that process reliabilism offers the best available solution to the puzzle of immediate justification. If this thesis can be made good, it will be a feather in reliabilism's cap, even if it doesn't solve other outstanding problems for the theory.

2. Feldman's Proposal: A Proper Response to Experience

Feldman, in his book Epistemology (2003), discusses the problem of justified basic beliefs with special attention to modest foundationalism. On Feldman's construal, modest foundationalism includes the thesis that immediate justification in the perceptual arena pertains to (selected) perceptual beliefs whose contents are external world propositions, such as "There is a tree on the hill." Feldman introduces the notion of a spontaneously formed belief, a belief that isn't formed by conscious inference from other beliefs. He then considers principle P1 as a possible principle for basic beliefs:³

(P1) Being spontaneously formed makes a belief (immediately) justified. (2003, p. 73)

This principle is quickly rejected, on grounds of excessive simplicity. To improve on P1, Feldman introduces the notion of a "proper response to experience." Only beliefs that are proper responses to experience provide (noninferential) justifiedness, Feldman implies. Here is how he explains the notion:

When you walk into a room, see a table, and form a belief that there is a table there, ... [w]hat seems central is that your belief is a proper response to the perceptual stimulus you have. It is a suitable thing to believe given that experience. To believe something that does not fit that experience at all, such as that there is an elephant in the room, would not be a proper response to that experience. To believe something that goes beyond what is revealed in the experience, such as that there is a table that is exactly 12 years old, would not be a proper response to that experience." (2003, p. 74)

Feldman then proceeds to use this notion in a revised principle for immediate justification:

(P2) A spontaneously formed belief is justified provided it is a proper response to experiences and is not defeated by other evidence the believer has. (2003, p. 74)

The adequacy of Feldman's solution to the problem of immediate justification obviously depends on what account is given of the newly coined phrase "proper response to experience". What makes a belief a proper or improper response to experience? Feldman gives further illustrations of positive and negative instances of this notion:

Other examples clarify the idea. Compare a novice bird-watcher and an expert walking together in the woods, seeking out the rare pink-spotted flycatcher. A bird flies by and each person spontaneously forms the belief that there is a pink-spotted flycatcher there. The expert knows this to be true but the novice is jumping to a conclusion out of excitement. The expert has a well-founded belief but the novice does not. In the same situation, both the novice and the expert may have well-founded beliefs about the color, shape and size of the bird they see. This suggests that there is some relevant difference between such properties as

being gray with pink spots and about 4 inches long and properties such as being a pink-spotted flycatcher. One might say that the former are “closer to experience” than the latter. Anyone with proper vision can discern the former in experience. This is not true of the latter. (2003, p.75)

These are instructive examples, but what general principles should be extracted from them? Feldman proposes two factors that seem relevant to being a proper response to experience.

First, when the contents of the belief are closer to the direct contents of experience, they are more apt to be properly based on experience. Second, modest foundationalism can say that training and experience affect what counts as a proper response to experience. The expert’s training makes her response proper. For beliefs that are more distant from experience, such training is necessary for the belief to be properly based on experience. (2003, p. 75)

Consider Feldman’s first factor for proper responsiveness, viz., closeness of the belief content to the content of experience. Feldman is obviously assuming that experience (e.g., visual experience) has propositional content. He assumes, in other words, that at least some immediate justifiers are contentful states. Some epistemologists would dispute this claim, of course, at least with respect to experiential justifiers. Even if the assumption is accepted, however, there are those who will argue that if a contentful state is to confer justification, *it itself* must be justified. If that were right, the justification it confers would no longer be *immediate*.

Setting this concern aside, Feldman’s appeal to the closeness-of-content factor presupposes that proper responsiveness is always a matter of a close match between the content of the experience and the content of the belief that responds to it. Is this right? Pryor (2005) doesn’t discuss Feldman directly, but his discussion brings out the problem of heavy appeal to content-closeness. Assume *arguendo* that some mental states don’t have propositional content, for example, pain states in general and headaches in particular. Suppose I am currently feeling pain, or having a headache, and I proceed to believe “I am now in pain” or “I have a headache.” Each of these beliefs is an excellent candidate for a justified basic belief. Or, formulating the matter in terms of propositional justification, I am justified in believing either of these propositions. Now it appears that the justification-conferring state is the state of being in pain or the state of having a headache, respectively. By hypothesis, however, neither of these justification-conferring states has content. So neither has any content that is *close* to the content of the proposition or belief in question. Content closeness is of no help here.

This problem is readily generalized to contentful mental states, such as desires. Ralph has an (occurrent) desire that the year 2008 be free of natural disasters. In an act of self-reflection (or introspection), Ralph forms the belief that he is currently experiencing *a desire*. This is a clear, traditional candidate for an immediately justified belief. But, again, the content-closeness factor doesn’t help rationalize its status as a basic belief. The belief’s content – *I currently experience a desire* -- has very little

proximity to the desire's content to which the belief is properly responsive, viz., *The year 2008 should be free of natural disasters*. Although the desire is contentful, the belief in question doesn't address its content. The believed proposition simply says that a specified *type* of mental state is currently instantiated in him.⁴

Does Feldman's second factor for proper responsiveness, i.e., suitable training, offer much help? The problem here is under-explanation. Feldman simply doesn't tell us (A) what constitutes suitable training with respect to a given belief, or (B) why suitable training should help confer immediate justifiedness. Presumably, suitable training does not consist in acquiring justified beliefs about the subject matter in question from which the target belief can be inferred. Because this would threaten to make the target belief indirectly rather than directly justified. What, then, does suitable training consist in? The most obvious answer is that suitable training makes beliefs of the appropriate category come out *true*. But this seems like a thinly disguised way of saying that suitable training breeds *reliability*; yet reliability is precisely what an internalist like Feldman is emphatically unwilling to accept as a ground or basis of justifiedness. If reliability is rejected, however, what is Feldman's story of why suitable training confers justifiedness?⁵

3. Pryor on Immediate Justification

Pryor (2005) has a very instructive discussion of immediate justification. But his discussion is directed squarely at the topic announced in his title, "There Is Immediate Justification." He is concerned with supporting the thesis *that there is* immediate justification, as contrasted with explaining *wherein* immediate justification consists. In other words, Pryor's discussion focuses on questions Q1 and Q2 rather than question Q3. He focuses on principles and arguments for the negative thesis that there is no immediate justification, and pokes holes in these arguments. He does an excellent job of expounding and exposing lacunae in such principles and arguments, but this doesn't provide a positive story of immediate justification. It doesn't help answer question Q3. Elsewhere Pryor does sketch a positive account of immediate justification (Pryor 2000). That account, however, closely resembles the more fully developed theory of Huemer, which will be discussed in section 4. The criticisms to be lodged against Huemer's view apply equally to Pryor's earlier proposal.

Let us examine a few of the principles and arguments against immediate justification examined and exposed in Pryor (2005). The central principle he examines is called "The Premise Principle":

The only things that can justify a belief that P are other states that assertively represent propositions, and those propositions have to be ones that *could be used as premises* in an argument for P. They have to stand in some kind of inferential relation to P; they have to imply it or inductively support it or something like that. (2005, p. 189)

Pryor diagnoses the chief argument behind the Premise Principle as follows. Consider a state without propositional content, like a headache. Since it has no propositional content, it can't stand in logical relations to the content of any beliefs. So why should it justify any one belief as opposed to others? For example, why should it justify the belief *I have a headache* as opposed to *I don't have a headache*, or any other belief? What can the defender of immediate justification (the foundationalist) say to make the justifying relations he postulates non-arbitrary? Pryor concedes that epistemologists must give principled, non-arbitrary rationales for the justifying relations they postulate. But he says, quite rightly, that it isn't obvious that they have to appeal to (states with) propositional contents to do it.

The second argument Pryor considers runs as follows:

[W]e ordinarily understand "justifications" for a belief to be *arguments* that support the belief. If you have *reasons* for your belief, they should be considerations you could in principle *cite*, or *give*, to someone who doubted or challenged the belief. You can't give someone else a non-propositional state like a headache (at least, not in the relevant sense); you can only give them *premises* and *arguments* that support your belief. This seems to show that justification and reasons *are* limited to things permitted by the Premise Principle. (2005, pp. 193-194)

Pryor responds, quite appropriately, that this argument appeals to the *dialectical* notion of a reason. This notion is different from the notion of a justification-maker, which is crucial for our inquiry. Pryor next distinguishes between two construals of the verb "justify". On the first construal, "justifying" a belief in P is a matter of proving or showing the belief to be just (or reasonable or credible). On the second construal, to "justify" something is akin to beautifying or electrifying it. When something beautifies a room, it does not *prove* that the room is beautiful; rather, it *makes* the room beautiful. Similarly, on the second construal, to justify a belief is a matter of *making* the belief just or reasonable, not a matter of *showing* the belief to be just (2005, p. 194). I think this is exactly right. However, it doesn't tell us what does make beliefs justified, especially in the case of directly or immediately justified beliefs.

Pryor provides one hint, however. He begins by distinguishing between there *being* reasons to believe P and one's *having* reasons to believe P. If *there are* reasons to believe P but you are unaware of them, they aren't reasons you have. And for anyone with internalist sympathies, such reasons cannot justify you in believing P. Thus, if something justifies you, it must be in some sense *available* to you. Although the availability constraint is attractive to internalists, externalists have offered reasons for doubting it (see Goldman, 1999b). Moreover, even if the availability constraint is accepted, such a constraint is only a necessary condition for justifiedness. And it gives us no inkling of *which* available states, facts, or conditions make *which* propositions or beliefs immediately justified. I am assuming here that we have been told -- which we haven't -- what exactly it means to be "available." Even knowing the appropriate

meaning of this term, however, wouldn't thereby give us an answer to question Q3, viz., which available states make which beliefs or propositions justified.

4. Huemer on Direct Perceptual Justifiedness

Huemer (2001) offers a general account of immediate justification, an account that is stated quite compactly:

(P3) If it seems to S as if P, then S thereby has at least prima facie justification for believing that P. (2001, p. 99)

Huemer calls this the rule of "phenomenal conservatism." The kinds of seemings (or appearances) he has in mind are of at least three sorts: perceptual seemings, memory-related seemings, and intellectual seemings (intuitions). Each gives rise, according to Huemer, to a species of justificational foundations. Three points of clarification should be noted. First, the principle of phenomenal conservatism concerns only prima facie, not ultima facie, justification. This is no different from the other approaches to immediate justification we have discussed, which are equally aimed at prima facie justification. Second, principle P3, as written, doesn't explicitly say anything about immediate justification. But since the sufficient condition it offers makes no requirement concerning prior beliefs (whether justified or not), it is implicitly a principle of immediate justification. Third, Huemer specifically claims that immediate, or foundational, justification of *all* kinds is conferred by seeming-as-if states:

I intend phenomenal conservatism to be a general principle of foundational justification. It explains not only why perceptual beliefs are noninferentially justified, but also why any other belief that is noninferentially justified is such.... I propose to account for perceptual knowledge by the same general principle I apply to all other kinds of knowledge. (2001, p. 102)

How shall we assess this theory? Markie (2005) lodges several criticisms of it. He begins by critiquing some arguments Huemer offers in defense of principle (3). I shall not review these arguments or Markie's criticisms of them, because even if Huemer's arguments for principle P3 are inadequate (and I definitely lean in this direction), the principle might still be correct. So let me turn to Markie's counterexamples against P3.

Markie says that mental processes that are incapable of producing prima facie justification can nonetheless determine how things seem to us. He gives an example to illustrate his point that is very similar to Feldman's bird-watchers example:

Suppose that we are prospecting for gold. You have learned to identify a gold nugget on sight but I have no such knowledge. As the water washes out of my pan, we both look at a pebble, which is in fact a gold nugget. My desire to discover gold makes it seem to me as if the pebble is gold; your learned identification skills make it seem that way to you. According to (PC) [our

principle (3)], the belief that it is gold has *prima facie* justification for both of us. Yet, certainly, my wishful thinking should not gain my perceptual belief the same positive epistemic status of defeasible justification as your learned identification skills. (2005, pp. 356-357)

Next Markie offers an example very similar to Feldman's contrast between believing there's a table present and believing there's a table exactly 12 years old present (both on the basis of the table's mere appearance). Markie's example and discussion run as follows:

Suppose that I perceive the walnut tree in my yard, and, having learned to identify walnut trees visually, it seems to me that it is a walnut tree. The same phenomenological experience that makes it seem to me that the tree is a walnut also makes it seem to me that it was planted on April 24, 1914. Nothing in the phenomenological experience or my identification skills supports things seeming this way to me. There is no date-of-planting sign on the tree, for example.... It is plausible that my experience directly justifies my belief that the tree is a walnut. The phenomenological character of the experience and my identification skills combine to support that belief. However, the phenomenological character of my experience does not combine with anything I have learned, or could learn, to support my belief that the tree was planted April 24, 1914.... My perception cannot directly justify my belief about the planting date. Nonetheless, according to (PC) [our principle P3], both my belief that it is a walnut tree and my belief that it was planted on April 24, 1914 are *prima facie*, and so defeasibly, justified for me. (2005, p. 357)

I agree with Markie that these cases are difficulties for Huemer's account. But the difficulties don't end there. Another problem with the account is that it doesn't cover all cases of immediate, or foundational, justification, as Huemer claims it does. In fact, the most classical type of direct justification isn't covered by phenomenal conservatism.

Suppose someone both experiences pain and thinks about Vienna, and suppose he is aware of each experience and therefore forms a belief that he is in pain and a belief that he is thinking about Vienna. These are classical cases of immediately justified beliefs. How does phenomenal conservatism handle them? To sustain phenomenal conservatism, there would have to be introspective seemings or appearances analogous to perceptual seemings and memory-related seemings. In addition to feeling pain and believing one is in pain, there would have to be a "seeming as if" one is in pain; and in addition to thinking about Vienna and believing that one is thinking about Vienna, there would have to be a "seeming as if" one is thinking about Vienna. But there aren't any such introspective seemings or appearances, distinct from the feelings and thought episodes themselves and the beliefs about them. So, introspective foundational beliefs constitute counterexamples to Huemer's phenomenal conservatism. Huemer doesn't mention such cases, despite the fact that they are core examples of basic beliefs according to epistemological tradition.

Huemer might reply to this objection as follows. In the special case of introspective beliefs, the mental state itself is the relevant appearance or seeming. A *distinct* seeming, however, isn't needed because the conscious state itself plays that role. This reply is unsatisfactory. The seemings Huemer posits for perception, memory, and a priori intellection are all mental states (or events) that resemble beliefs in two respects. First, they have propositional content: they are of the form "it seems as if P." Second, seemings are descriptive or depictive attitudes; they have a mind-to-world direction of fit rather than a world-to-mind direction of fit. Not all mental states or episodes, however, have these properties. States like being in pain (arguably) have no propositional content at all. And states like desires, wishes, and intentions are not descriptive or depictive attitudes; they have a world-to-mind direction of fit rather than a mind-to-world direction of fit. Nonetheless, all these kinds of mental states are just as capable of conferring, or helping to confer, immediate justification as states like thinking that P. One can be immediately justified in believing that one is in pain, and one can be immediately justified in believing one has a desire, or desires that Q. However strenuously Huemer may work to make introspection cases fit the "seeming" mold, it won't work for all cases. In fact, it is plausible *at most* for a small subset of cases. For the rest, a different story is needed. This casts doubt on the approach in general.

4. Process Reliabilism and Perceptual Foundations

Let me turn now to some positive proposals of my own. As indicated in the introduction, these proposals will have a process-reliabilist theme. There is wide consensus that if foundationalism is correct, there are probably multiple domains of basic beliefs or propositions. This is clearly illustrated by Huemer, who presents three categories of basic beliefs: perceptual, memory-related, and intellectual (intuitive). This section deals with the first of these categories: perceptual basic beliefs.

Parenthetically, we should say a word or two about the relationship between reliabilism and foundationalism. My original formulation of process reliabilism (Goldman 1979) proposed a recursive format for justified beliefs. In this format, base clauses state sufficient conditions for justifiedness without invoking the justificational status of other beliefs. Recursive clauses state sufficient conditions for justifiedness that do invoke the justificational status of other beliefs. Obviously, base clauses would permit immediately, directly, or foundationally justified beliefs. Sosa (1980) calls this "formal foundationalism." Formal foundationalism does not entail process reliabilism, however. It is not committed to any role for causal processes in the conditions for justifiedness, or to reliability as a crucial determinant of justificational status. Thus, formal foundationalism is compatible with process reliabilism, but process reliabilism is only one species of formal foundationalism. Traditional foundationalism, which spurned causal factors as irrelevant to justificational status, would be another variety of formal foundationalism.

I turn now to the topic of perceptual foundational beliefs. This is a class of perceptual beliefs about the external world that allegedly qualify as immediately (*prima facie*) justified. Not all foundationalists, of course, countenance perceptual foundational

beliefs. Cartesian foundationalism restricts foundational beliefs to beliefs about one's own current mental states (and elementary propositions of reason). But modest foundationalists typically admit the existence of basic beliefs about the external world. The present question is: What conditions, exactly, must these beliefs fulfill to be directly justified?

Let us pursue the clues provided by Feldman's examples. Although his proper-response-to-experience theory is vague and inadequate, his examples are helpful. Markie's examples are cast in the same mold and are equally instructive, but I'll stick with Feldman's.

Consider the difference between Feldman's expert and novice bird watchers. Assume for the sake of discussion that the two have identical visual experiences. How do they differ so that the expert's belief in the species-identifying proposition ("It's a pink-spotted flycatcher") is immediately justified and the novice's belief in the same proposition isn't so justified? The difference might well lie in differences in the cognitive processes by which they respectively proceed from visual experience to belief.⁶ The expert presumably connects selected features of his visual experience to things he has stored in memory about pink-spotted flycatchers, perhaps securing an appropriate "match" between features in the experience and features in the memory store. The novice does no such thing; he just guesses (according to Feldman's story). Presumably, the token process used by the expert is an instance of a *generally reliable* (truth-conducive) kind of cognitive process. It involves classifying an instance of a category-type by matching observed features of the instance to information about the category. This is a much more truth-conducive classification process -- as befits an expert -- than mere guessing. Thus, Feldman's vague notion of a *proper* response to experience might be spelled out as a doxastic response to experience that is the output of a reliable belief-forming process applied to that experience (as input). Since all of the information processing that occurs in the expert's cognitive system is assumed here to be unconscious, we don't violate the requirement that the belief be formed spontaneously or noninferentially; so it remains a candidate for being immediately justified.

Feldman's 12-year-old table example can be handled similarly. If a person sees a table for the first time, there won't be any clues in his visual experience of the table to which a *reliable* belief-forming process could be applied that would generate the output that the table is 12 years old. Any process that might be used to generate this belief from his visual experience (without any relevant background information about the table) would not be reliably formed. By contrast, it would be easy to have a reliable belief-forming process that could be applied to the same visual experience and generate the output that a table is present. Where Feldman speaks of the two beliefs as being, respectively, an "improper" and a "proper" response to the visual experience, the difference might be captured more precisely by saying that one belief stems from applying an unreliable process to the visual experience (or by applying the process without engaging the visual experience at all) whereas the other stems from an application of a reliable process to the visual experience.

I don't attempt here to give a tight formulation of "reliable." For example, I don't try to specify the domain (modal or otherwise) in which the truth-ratio for a process is to be determined, or to say how processes are to be typed (the generality problem). These are questions for the general theory of reliabilism, which I intend to skirt here.⁷ My aim is only to display the promise of process reliabilism as a general approach to the core foundationalist idea of immediate justification. The aim is to show, or at least illustrate, why it is more promising than competing approaches. I am not trying to provide all the details.

In characterizing the causal process relevant to the justifiedness of a perceptual belief, I have thus far followed Feldman in considering the process *from experience to belief*. This accords with his suggestion that perceptual justification (primarily) consists of making a proper doxastic response *to experience*. An alternative treatment, however, is available to process reliabilism. The relevant process might be a more extended one, not *from experience to belief* but *from receptor-stimulation (to experience) to belief*. This more extended process has potential advantages because of what might transpire between receptor stimulation and experience. It might include events or operations influenced by prior training, and these types of events might contribute to reliability. Thus, we should leave the door open to either type of process – the more extended or the less extended process -- as the optimal choice from the perspective of process reliabilism. However, I won't attempt to make a firm resolution of the issue.⁸

How would process reliabilism deal with *propositional* (as opposed to *doxastic*) justification for perceptual propositions? Suppose that in Feldman's bird-watcher example neither the expert nor the novice actually believes that the bird is a pink-spotted flycatcher, although they both have a visual experience that results from seeing it. We want to distinguish between them by saying that the expert, but not the novice, is propositionally justified in having this belief. How would reliabilism explain the difference? Where direct justification is in question, process reliabilism can say the following. A proposition P is propositionally justified for an epistemic agent S (at time t) just in case S's total mental state (at t) is such that *if S were to "target" proposition P and were to apply suitable reliable processes in her repertoire to that total state, then a belief that P would be generated in S* (see Goldman 1979, p. 21, or, in the reprinted version, Goldman 1992, p. 124). In the bird-watcher example, the expert's total state and the reliable processes in her repertoire are such that if she targeted the pink-spotted flycatcher proposition and applied (a suitable subset of) her reliable processes to her total state at the time, then a belief would be generated that the sighted bird is a pink-spotted flycatcher. By contrast, the novice does not satisfy these conditions. There are no reliable processes in her repertoire such that, if they were applied to her total state, they would generate the pink-spotted flycatcher belief. Of course, the combination of her total state and *unreliable* processes in her repertoire are capable of generating this belief (that is how matters transpire in Feldman's original version of the case); but that doesn't meet the specified condition. So process reliabilism can account for the difference between the expert and the novice in both the doxastic and the propositional senses of justification.

5. Process Reliabilism and Introspective Foundations

Although contemporary foundationalists often allow basic beliefs to include perceptual beliefs with external-world contents, they also typically allow the set of basic beliefs to include beliefs with first-person mental-state contents, e.g., *I am in pain*. We used this point in critiquing Huemer's account. In this section we ask how process reliabilism hopes to handle this subclass of basic beliefs. One question is: What cognitive process would reliabilism invoke? A second question to be addressed is how process reliabilism would handle both the propositional and the doxastic interpretations of justification in this domain.

Starting with the first question, the obvious process of interest is the process of introspection, or inner sense – assuming there is such a process. Elsewhere I endorse the existence of such a process as the best explanation of people's first-person mental-classification activity (Goldman 2006, chap. 9). Alternative philosophical and psychological accounts of mental self-classification are surveyed and found wanting. The introspective process must be rather complex, however. I suggest that it involves two stages: first, a stage of *attending* to selected portions of the mental field, and second, a stage of *classifying* what is found there. A classificational act is a judgment or belief; so, in its second stage, introspection is a species of belief-forming process.

Suppose introspective classification is a generally reliable process (not to say infallible). Then any belief formed by introspection, according to process reliabilism, will be justified. Will it be immediately, or directly, justified? If it is spontaneous, i.e., not consciously inferential, then it should be considered immediately justified. And this seems overwhelmingly plausible. When we introspectively classify our mental states, there isn't any conscious inference from other beliefs. True, we may unconsciously draw on information about what it is like (from an internal perspective) to have a pain, a desire, or a belief. From a psychological point of view, however, this need be no different from cases of vision-based classification, where we draw unconsciously on information about what trees or tables look like. If visual object classification is allowed to yield immediately justified beliefs, the case for introspective mental-state classification is at least as good. However, as noted in section 3, immediate justification based on introspection does not depend on any introspective "seeming" or "appearance" on which the belief is based.

The foregoing discussion addresses the doxastic interpretation of justification. What about the propositional interpretation? Can we make sense in process-reliabilist terms of being potentially justified – of being in a position to be justified -- in believing a proposition about a current mental state, even when one doesn't believe it? Yes, process reliabilism has a fairly straightforward story to tell. If one feels pain at a given moment, one can apply the first stage of an introspection operation to that feeling, i.e., one can *attend* to it, and then proceed to the second stage of classification. If one performed both of these operations, then the belief that one feels pain would be generated, and this would be the product of a reliable classificational operation. So this is the sense in which one is in a "position" to be justified, or is potentially justified, in believing one is in pain even when one doesn't actually hold that belief.

6. Immediate Justification and the Internalism/Externalism Dispute

The usual ways of describing foundationalism and its key notion of immediate justification are neutral vis-à-vis the internalism/externalism controversy. In advancing process reliabilism as the best detailed account of immediate justification, however, I am promoting the cause of externalism. How might internalists respond? As a staunch internalist, Feldman obviously tries to provide an internalist account of direct perceptual justification. His key notion of “proper response to experience” is evidently crafted so as to omit all unpalatable externalist ingredients. By contrast, I argue that the only satisfactory way to flesh out the notion of “proper response” is to utilize the notion of reliability, a paradigmatically externalist concept.

Elsewhere Conee and Feldman (2001) specify the general kind of internalism about justification they mean to defend. They don't address immediate justification in particular, but their general approach to justification is clearly intended to apply to this variety of justification. They first distinguish two different conceptions of internalism, viz., accessibilism and mentalism. Accessibilism holds that the justification of a person's beliefs is determined by things to which the person has special access, e.g., “direct” access, or access through “introspection or reflection.” Mentalism holds that the justificational status of beliefs is fixed only by things internal to the cognizer's mental life.

Conee and Feldman favor the mentalistic brand of internalism. They say that internalism, so construed, is committed to two theses. The first thesis, (S), asserts the strong supervenience of epistemic justification on the mental:

(S) The justificatory status of a person's doxastic attitudes strongly supervenes on the person's occurrent and dispositional mental states, events and conditions. (2001, p. 234)

The second thesis spells out a principal implication of (S):

(M) If any two possible individuals are exactly alike mentally, then they are alike justificationally, e.g., the same beliefs are justified for them to the same extent. (2001, p. 234)

Suppose we accept these two theses. Do they provide the resources for resolving the justificational status of beliefs in the kinds of cases of interest here, for example, Feldman's own cases in *Epistemology*? No. Principle (M) provides a single sufficient condition for the *sameness* of justificational status, but it doesn't provide conditions, either necessary or sufficient, for fixing *positive* or *negative* justificational status. Hence, it doesn't meet the challenge of question Q3.

To illustrate, consider the two bird watchers, who both form a spontaneous belief that the sighted bird is a pink-spotted flycatcher. In Feldman's description, the novice and expert are not exactly alike mentally (even with respect to this belief). So their two

beliefs are allowed to differ in justificational status. Does either (M) or (S) provide materials for determining that the expert's belief *is* justified and that the novice's belief *isn't* justified? No. The theoretical resources of the approach, as specified by (M) and (S), simply aren't adequate to draw any such conclusion, even when we add the psychological details Feldman supplies (namely, that the expert has lots of relevant training, that the novice jumps to a conclusion out of excitement, etc.). Of course, Feldman tries to provide an additional theoretical principle when he talks about "proper responses to experience," and tries to explain this in terms of beliefs having varying degrees of "closeness to experience." Greater closeness is supposed to generate greater justifiedness; and the closeness motif is arguably a matter of mentality, and hence (arguably) a legitimate factor for internalism to appeal to.⁹ As discussed earlier, however, the "closeness" idea doesn't really work. What else might internalism have on offer? I know of no other promising internalist idea.

On the other hand, appeal to process reliability generates answers that precisely match our intuitions about cases. So this externalist ingredient performs exactly the job we want performed. Reliabilist externalism has an exemplary solution to our problem, whereas mentalistic internalism (thus far) has no satisfactory solution at all. To be sure, the proposed reliabilist solution to the problem of immediate justifiedness doesn't address any of the outstanding problems for reliabilism. But insofar as reliabilism seems to have a far better solution to the problem of immediate justifiedness than any of its rivals, it receives an extra infusion of support.¹⁰

References

- Alston, W. (1995). How to think about reliability. *Philosophical Topics* 23: 1-29.
- Beebe, J. (2004). The generality problem, statistical relevance, and the tri-level hypothesis. *Nous* 38: 177-194.
- Chisholm, R. (1977). Theory of Knowledge, 2nd edition. Englewood-Cliffs, NJ: Prentice-Hall.
- Comesana, J. (2005). We Are (Almost) All Externalists Now. In J. Hawthorne, ed., *Philosophical Perspectives* 19, *Epistemology*. Malden, MA: Blackwell.
- Conee, E. and Feldman, R. (2001). Internalism defended. In H. Kornblith, ed., *Epistemology: Internalism and Externalism*. Malden, MA: Blackwell.
- Feldman, R. (2003). *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Goldman, A. (1979). What is justified belief? In G. Pappas, ed., *Justification and Knowledge*. Dordrecht: Reidel. Reprinted in A. Goldman, *Liaisons: Philosophy Meets the Cognitive and Social Sciences* (1992). Cambridge, MA: MIT Press.
- Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. (1988). Strong and weak justification. In J. Tomberlin, ed., *Philosophical Perspectives* 2, *Epistemology*. Atascadero, CA: Ridgeview. Reprinted in A. Goldman, *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press.
- Goldman, A. (1992). Epistemic folkways and scientific epistemology. In *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge, MA: MIT Press.
- Goldman, A. (1999a). A priori warrant and naturalistic epistemology. In J. Tomberlin, ed., *Philosophical Perspectives* 13, *Epistemology*. Malden, MA: Blackwell. Reprinted in A. Goldman, *Pathways to Knowledge, Private and Public* (2002). New York: Oxford University Press.
- Goldman, A. (1999b). Internalism exposed. *Journal of Philosophy* 96: 271-293. Reprinted in A. Goldman, *Pathways to Knowledge: Private and Public* (2002). New York: Oxford University Press.
- Goldman, A. (2001). Replies to the contributors. *Philosophical Topics* 29: 461-511.

- Goldman, A. (2006). *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. New York: Oxford University Press.
- Greco, J. (2000). *Putting Sceptics in Their Place*. Cambridge: Cambridge University Press.
- Heller, M. (1995). The simple solution to the generality problem. *Nous* 29: 501-515.
- Huemer, M. (2001). *Skepticism and the Veil of Perception*. Lanham, MD: Rowman and Littlefield.
- Markie, P. (2005). The mystery of direct perceptual justification. *Philosophical Studies* 126: 347-373.
- Pryor, J. (2000). The skeptic and the dogmatist. *Nous* 34: 517-549.
- Pryor, J. (2005). There is immediate justification. In M. Steup and E. Sosa, eds., *Contemporary Debates in Epistemology*. Malden, MA: Blackwell.
- Sosa, E. (1980). The raft and the pyramid: Coherence versus foundations in the theory of knowledge. *Midwest Studies in Philosophy* 5, *Studies in Epistemology*. Minneapolis: University of Minnesota Press.
- Sosa, E. (1991). *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Sosa, E. (2001). Goldman's reliabilism and virtue epistemology. *Philosophical Topics* 29: 383-400.

¹ For example, Roderick Chisholm (1977) proposed to think of the “directly evident” as that which “constitutes its own evidence”, and characterized it in terms of states that are “apprehended through themselves” (1977, p. 25).

² I have proposed various reformulations of process reliabilism in a number of works that followed my initial formulation of the theory in Goldman 1979. See especially Goldman 1986, 1988, 1992, 1999a, 2001. These reformulations were intended to answer or resolve several problems, including the generality problem. Many other writers have also proposed reformulations or solutions, including Alston 1995, Beebe 2004, Greco 2000, Heller 1995, and Sosa 1991, 2001. Further efforts at improving the formulation of reliabilism are underway, but won’t be explored here.

³ All of the numbered principles are relabeled from their original locus of publication.

⁴ An anonymous referee suggests that Ralph’s belief isn’t a response to *his desire* that the year 2008 be free of natural disasters. Rather, it is a response to *his being aware of* his desire that the year 2008 be free of natural disasters. And the content of this awareness is very close indeed to the content of the belief. One reply to this suggestion is that it isn’t clear what content is possessed by a *de re* awareness (an awareness of another state). Setting this point aside, why should we posit the existence of such an awareness? Why should we suppose that, in addition to having a desire that the year 2008 be free of natural disasters and a belief concerning that desire, Ralph also has a separate and distinct awareness of this desire? When Feldman discusses a belief that properly responds to a perceptual experience, he doesn’t posit an extra mental state – an awareness of the perceptual experience – that is layered *in between* the target mental state and the belief. Why make such a posit here? To be sure, in the perceptual-experience case the belief is properly responsive to the content of the experience. But why shouldn’t a belief be equally capable of being properly responsive to a mental state’s *type* as well as to its content? Intuitively, proper responsiveness to a state’s type seems as good a candidate for grounding immediate justification as proper responsiveness to its content. But, of course, this doesn’t fit with Feldman’s characterization of proper responsiveness.

⁵ Feldman himself concedes that he has not adequately explained the notion of a proper response to experience. He remarks that “a more fully developed account of the conditions under which a belief is properly based on experience is desirable” (2003, p. 78).

⁶ Interestingly, Markie explicitly adverts to cognitive processes in describing his versions of this kind of case. I think he is right to do so, because it makes the cases understandable. However, although he has latched onto this component of what I regard as the right solution to the problem, his dismissal of reliabilism at the beginning of his paper keeps him from seeing the full (reliabilist) solution. So his article’s theme is that the problem remains a “mystery”.

⁷ I do not doubt, however, that answering these questions could help assuage some worries about the reliabilist approach. Consider, for example, the worry that merely “accidental” reliability isn’t sufficient for immediate justification. (Rich Feldman expressed this worry in a personal communication.) This problem might be partly resolved by selecting an appropriate domain for fixing the truth-ratio of a process. Requiring a process to have a high truth-ratio in nearby possible worlds, as well as in the actual one, might help exclude processes that are only accidentally reliable. An additional constraint might be added concerning the source, or origin, of a process’s reliability. I discuss this kind of constraint in *Epistemology and Cognition* (Goldman 1986, pp. 51-53, 94-95).

⁸ I am grateful to Jim Pryor for emphasizing the fact that there is a choice to be made about process extent. He personally recommends the choice of an even more extended process than either of the two considered in the text, namely, a process that reaches from *external stimuli* to experience. (At least he recommends this for the analysis of *propositional* justification, which is addressed in the next paragraph of the text.) However, process reliabilism was formulated from the start as a theory that highlights *cognitive* processes,

and cognitive processes are best interpreted as operations *internal* to the organism. See Goldman 1979, pp. 12-13; or, in the reprinted version, Goldman 1992, p. 116.

⁹ Even this assumption is dubious, however. Content closeness seems to be either a species of evidential support or in the same family as the support relation. But the support relation is not rightly considered a mental state of affairs, or relation, as argued in Goldman 1999b and Comesana 2005. If this is correct, then any notion like closeness is not really congenial to mentalistic internalism.

¹⁰ Thanks to Rich Feldman, Jim Pryor, and two anonymous referees for helpful comments on previous versions of this paper.