

Why the Generality Problem is Everybody's Problem

Michael A Bishop
Department of Philosophy
Florida State University
356 C Dodd Hall
Tallahassee, FL 32306
mbishop@fsu.edu
Phone: 850 644-4127
Fax: 850 644-3832

ABSTRACT. The generality problem is widely considered to be a devastating objection to reliabilist theories of justification. My goal in this paper is to argue that a version of the generality problem applies to all plausible theories of justification. Assume that any plausible theory must allow for the possibility of reflective justification – S's belief, B, is justified on the basis of S's knowledge that she arrived at B as a result of a highly (but not perfectly) reliable way of reasoning, R. The generality problem applies to all cases of reflective justification: Given that B is the product of a process-token that is an instance of indefinitely many belief-forming process-types (or BFPTs), why is the reliability of R, rather than the reliability of one of the indefinitely many other BFPTs, relevant to B's justificatory status? This form of the generality problem is restricted because it applies only to cases of reflective justification. But unless it is solved, the generality problem haunts all plausible theories of justification, not just reliabilist ones.

Why the Generality Problem is Everybody's Problem

My thesis is that some version of the generality problem afflicts all plausible theories of justification.¹ This is bad news because the generality problem is widely thought to be devastating. In fact, Conee and Feldman (1998) argue that theories that are subject to the generality problem “look hopeless” (24) and that the prospects for a solution to the problem “are worse than bleak” (5). After introducing the generality problem (section 1), I note that any plausible view of justification must allow for cases of reflective justification – S’s belief is justified on the basis of S’s knowledge that she arrived at it as a result of a highly (but not perfectly) reliable way of reasoning (section 2). I will then argue that if the generality problem has no solution, no theory of justification can account for reflective justification; and so the generality problem bedevils all plausible theories of justification (section 3). One might think that in cases of reflective justification the generality problem is relatively easy to solve. I consider four prima facie attractive but doomed solutions and argue that there is no reason to suppose that the restricted form of the problem (which applies only to cases of reflective justification) is any easier to solve than the familiar, unrestricted form of the problem (section 4). So we are faced with a dilemma: Either the generality problem has a solution or there is no plausible theory of justification.

¹ Others have made similar points. For example, Heller (1995) and Adler & Levin (2002) argue that the generality problem arises for more than just reliabilism. Comesaña (2006) argues that “the problem with the characterization of well-foundedness is intuitively the same as the generality problem, and so any theory that incorporates the notion of well-foundedness into the characterization of justification will have to deal with the generality problem” (31). The argument I propose here, while similar in spirit to the arguments proposed by these authors, is different in detail and broader in scope.

1. The generality problem

Reliabilists hold that a newly minted belief's justificatory status depends on the reliability of the belief-forming process that produced it. But a token belief-forming process is either perfectly reliable (if the resulting belief is true) or perfectly unreliable (if it is false). Since there are true unjustified beliefs and false justified beliefs, the process relevant for determining the justificatory status of the belief cannot be the token that produced it (Goldman 1979, but see Comesaña 2006). And here is the challenge posed by the generality problem: The reliabilist must *categorize* process-tokens into process-types for the purposes of epistemic evaluation. This is a particularly tough problem for two reasons. First, there are indefinitely many ways to categorize belief-forming process-tokens into types. As Conee and Feldman note, "[t]he number of types is unlimited. They are as numerous as the properties had by the belief-forming process" token (1998, 2). This assumes a very permissive view of what counts as a belief-forming process-type (BFPT) – any set of belief-forming process-tokens that share some property or properties. While some have criticized this permissiveness (e.g., Wallis 1994, Alston 1995), I will embrace it in arguing that the generality problem applies to all plausible theories of justification. The second reason the generality problem is hard is that the reliabilist must propose a *principled* way to identify the BFPT that is relevant to epistemological assessment. Feldman explains why this is essential:

The fact that every belief results from a process token that is an instance of many types, some reliable and some not, may partly account for the initial attraction of the reliability theory. In thinking about particular beliefs one can first decide intuitively whether the belief is justified and then go on to describe the process responsible for the belief in a way that appears to make the theory have the right

result. Similarly, of course, critics of the theory can describe processes in ways that seem to make the theory have false consequences (1985, 160).

Without a principled solution to the generality problem, the reliabilist can always start with the intuitively correct judgment about a belief's justificatory status and then cherry pick a BFPT that yields that judgment. A principled solution to the generality problem will be "universal... in that it must specify the relevant type whenever there are definite facts about justification" and it will not involve "*ad hoc* case-by-case selections of relevant types that match our intuitions" (Conee and Feldman 1998, 4).

For every possible belief that has a justificatory status, a solution to the generality problem must articulate a principled way to identify the one BFPT (out of indefinitely many potential BFPTs) whose reliability is relevant to determining that belief's justificatory status.² No wonder there is so much pessimism. Conee and Feldman (1985) argue that the prospects for a solution to the generality problem "are worse than bleak" (5) and that theories that are subject to it "look hopeless" (24). For their part, Pollock and Cruz "can see no way around" the generality problem (1999, 118). It's fair to say that many philosophers share this pessimism. But the pessimists have not appreciated that some version of the problem bedevils all plausible theories of justification. The generality problem rears its ugly head *any time* a belief's justificatory status is, or is taken to be, a function of the reliability of its production.³ If you take belief B to be justified because it was reliably produced, regardless of your epistemological persuasion, the challenge posed by the generality problem immediately kicks in: Given that B is the product of a

² The assumption that a solution to the generality problem must identify a single relevant BFPT has been denied by Charles Wallis (1994) and Mark Wunderlich (2003). I propose to grant it for the sake of argument.

³ What if someone takes a belief's justificatory status to be a function of the reliability of its production but is mistaken – it is a function of something else? It doesn't matter. The generality problem arises for that attempt to show that the belief is justified – even if that attempt is ultimately doomed.

process-token that is an instance of indefinitely many BFPTs, why is the reliability of *that* BFPT relevant to B's justificatory status rather than the reliability of one of the indefinitely many others? If *everyone* is committed to taking the justificatory status of *some* beliefs to be a function of the reliability of their production, then the generality problem threatens all plausible theories of justification.

2. Three cases of reflective justification⁴

My goal in this section is to describe three target beliefs, two justified and one unjustified. All three are preceded by reflective, second-order reasoning. The idea here is familiar. Before tackling a problem or a series of problems, we sometimes think strategically about how best to reason about that sort of problem. Of course, we can't always engage in this kind of reflective reasoning, as that would lead to infinite regress (i.e., third-order reflection about how best to reason about the second-order problem, etc.). But sometimes as a result of our second-order reflections, we come to believe that some way of reasoning is the best available way to think about a certain class of issues. And as a result, we might commit ourselves to a reasoning policy of the form: *Ceteris paribus, I will reason thus-and-so on problems of type P*. If all has gone well, when we follow that reasoning policy, we will come to a justified belief.

Case 1: Heart Attack

Sandra is a doctor who has to determine whether Sam, a heart attack patient, is high risk: Will he die of a heart attack in the next 30 days without special treatment? As it happens, Sandra knows that the following simple model (a decision tree) has been shown to be more reliable at making these sorts of judgments than doctors (Breiman et al. 1993, 2).

(1) Is the minimum systolic blood pressure over the initial 24 hour period > 91 ?

⁴ This is meant to echo Sosa's notion of reflective knowledge (1985).

NO: High Risk

YES: Is age > 62.5?

NO: Low Risk

YES: Is sinus tachycardia present?

YES: High Risk

NO: Low Risk

Sandra knows that this is the most reliable method available for making these sorts of judgments. As a result, Sandra commits herself to using this decision tree to make judgments about whether heart attack patients are high risk. Sandra's *first-order* evidence and beliefs are directly about Sam (e.g., his age, his blood pressure). Her *second-order* evidence and beliefs are about the different methods of forming beliefs on the basis of the first-order evidence. Suppose Sandra's first-order and second-order beliefs are all justified. No plausible epistemological theory is going to have trouble explaining this.⁵ So far, so good. But now, Sandra follows her policy and applies the above model to her first-order evidence about Sam, and she comes to believe that (say) Sam is low risk. What justifies this belief? The intuitively plausible suggestion is that Sandra's belief is justified *because she knows that the model is highly reliable and she knows that she used the model to arrive at that belief*. This justification appeals essentially to the BFPT that consists of the belief-forming process-tokens that share the property of instantiating the heart attack model. Of course, Sandra might acquire additional first-order evidence (e.g., that Sam is having arm and chest pains) that would justify revising her model-based belief. But absent such countervailing evidence, Sandra's first-order belief about Sam is justified.

⁵ Actually, I will take this back in the next section. Insofar as Sandra's beliefs about the reliability of the heart attack model are the result of testimony from experts she knows (or justifiably believes) are highly reliable, the generality problem arises for these beliefs as well.

Case #2: Reliable Testimony

Suppose on the basis of overwhelming evidence, you know that S is highly reliable about U.S. history, and you know that *ceteris paribus* the most reliable way for you to reason about U.S. history is to defer to S's judgments. As a result of this knowledge, you adopt a reasoning policy of (*ceteris paribus*) deferring to S about U.S. history. Now suppose that S tells you that Lincoln was the 16th President of the U.S., and as a result of your reasoning policy, you come to believe this. Surely, your belief is justified. And your belief about Lincoln is justified (in part) on the basis of your second-order knowledge about the reliability of a BFPT (consisting of the process tokens that involve you deferring to S's judgments about U.S. history).⁶

Case #3: Gerrymandered

After years of neglect, Paula spends much of the day, from noon to 4:55, balancing her checkbook, performing a few thousand simple arithmetic calculations. She spends the next five minutes coming to various absurdly optimistic beliefs about her future. She tries to justify these beliefs by appealing to a BFPT that she knows is reliable. She argues: "My beliefs about my future will be produced by process-tokens that are instances of a highly reliable BFPT. That BFPT consists of the process tokens that share the property of having produced my beliefs between noon and five."⁷ In fact, this BFPT is more reliable than the BFPTs appealed to in Heart Attack and Reliable Testimony. If those beliefs can be justified by appealing to reliable BFPTs, I

⁶ Testimonial knowledge has received considerable attention recently. (For an excellent overview of the literature, see Lackey 2006.) Perhaps the central issue concerning testimonial knowledge is whether testimonial beliefs are (in some sense) epistemically basic. I take no stand on that issue. My thesis requires only the uncontroversial claim that some testimonial beliefs can be justified by explicit, second-order knowledge that the testimony comes from a highly reliable source.

⁷ This argument takes advantage of Feldman and Conee's permissive view of BFPTs: A BFPT is any set of belief-forming process-tokens that share some property or properties.

can surely justify my belief by appealing to an even more reliable BFPT.” Of course, despite Paula’s appeal to an extremely reliable BFPT, her absurdly optimistic beliefs are not justified.

3. Why the generality problem is everybody’s problem

The target beliefs in Heart Attack and Reliable Testimony are clearly justified. But *if* the generality problem has no solution, then no theory of justification can yield the result that those beliefs are justified. The reason is that in both cases, the justification of the target belief must appeal essentially to the reliability of a BFPT. And this immediately kick-starts the generality problem: Of the indefinitely many BFPTs that have as an instance the process-token that produced the target belief, why is *that* BFPT relevant to the justificatory status of the target belief? If there is no solution to the generality problem, there is no principled answer. And without a principled answer, no theory of justification can account for why the target beliefs are justified.

To see this, let’s start with a reliabilist, Alvin, who argues that in Heart Attack Sandra’s belief that Sam is low risk is justified on the grounds that it was reliably produced. The pessimist about solving the generality problem argues: “Alvin, Sandra’s target belief was produced by a process-token that is an instance of indefinitely many BFPTs, some reliable, and some not. Without a solution to the generality problem, you have no legitimate reason to believe that *that* reliable BFPT is relevant to the justificatory status of the belief in question rather than one of the unreliable ones. Therefore, you cannot claim that the target belief is justified by the reliability of *that* BFPT.” And so the generality problem scuttles Alvin’s attempt to explain the justification of Sandra’s target belief. Alvin knows that the target belief was reliably produced. But if there is no solution to the generality problem, that knowledge cannot ground Alvin’s contention that the target belief is justified. But if no one has a solution to the generality problem, then why does

only Alvin get such rough treatment? Suppose the proponent of a different theory of justification – or the believer herself – knows that the target belief was reliably produced. How can that knowledge ground the contention that the target belief is justified? If there is no solution to the generality problem, it can't.

Now let's consider how an internalist, Richard, might handle Heart Attack. He argues that Sandra's target belief is justified because the rest of Sandra's mental states (especially her beliefs) bear the appropriate justification-conferring relations to the target belief. What are these justification-conferring mental states? Prima facie, they must include Sandra's second-order beliefs to the effect that her target belief was reliably produced. (Richard might try to deny this and thus avoid the generality problem. We will consider two such efforts in section 4.) The pessimist about solving the generality problem argues: "Richard, you contend that Sandra's second-order beliefs about how her target belief was reliably produced help confer justification on her target belief. But how can this be? After all, her target belief was produced by a belief-forming process-token that is an instance of indefinitely many BFPTs, some reliable and some not. There is no solution to the generality problem. So there is no legitimate reason to take *that* reliable BFPT, rather than some unreliable BFPT, to be relevant to the justificatory status of Sandra's target belief. Therefore, Sandra's second-order beliefs cannot confer justification on her target belief. Of course, those beliefs *could* confer justification on her target belief *if* the reliability of *that* BFPT were relevant to the justificatory status of her target belief. But there is no solution to the generality problem. So there is no reason to think that that BFPT is relevant to the justificatory status of Sandra's target belief." In support of this argument against Richard's attempt to account for the justificatory status of Sandra's target belief, the pessimist can appeal to Feldman and Conee: "Philosophers... purport to determine whether or not a given belief is

justified... using nothing more than one description of the process causing the belief. No such inference is acceptable” (1998, 3). If such an inference is unacceptable for reliabilist philosophers, it’s unacceptable for everyone else, too.

The same problem arises for Reliable Testimony. Your Lincoln belief was produced by a belief-forming process-token that is an instance of indefinitely many BFPTs, some reliable (e.g., defer to S about U.S. history), some not (e.g., defer to whoever is physically closest to you about U.S. history). What makes one of these BFPTs the one that is relevant to determining the justificatory status of your Lincoln belief? This is precisely the challenge set forth by the generality problem (Shafer-Landau 2003, 284). If there is no solution to it, then you have no reason to take the reliable BFPT – rather than one of the many unreliable BFPTs – to be relevant to the justificatory status of your Lincoln belief. And so the internalist cannot properly explain the justificatory status of your Lincoln belief. More precisely, the internalist cannot explain why your Lincoln belief is *more* justified than a run-of-the-mill testimonial belief – which it must be given that you know that the testifier is an expert on the topic.

These absurd consequences arise because the generality problem is so deeply corrosive. The case for the generality problem makes no essential appeal to reliabilism. If the generality problem cannot be solved, then for any belief and token process that produced it, *nobody* (reliabilist or not) has a principled way to identify the BFPT whose reliability is relevant to that belief’s justificatory status. And since some beliefs can only be justified by appeal to their reliable production, there had better be a solution to the generality problem.

4. Four failed solutions to the restricted form of the generality problem

For the reliabilist, the challenge posed by the generality problem is unrestricted – it applies to all beliefs. But for non-reliabilists, the challenge is restricted to beliefs that are

justified by appeal to reflective second-order reasoning, as in Heart Attack, Reliable Testimony, as well as scientific beliefs that are justified by virtue of having been produced by reliable methods (Adler & Levin 2002). Non-reliabilists might hope that the restricted form of the problem is much easier to solve than its unrestricted cousin. If this were true, we would be back to the status quo, with the generality problem dooming only reliabilism. My goal here will be to argue that there is no reason to suppose that the restricted form of the problem is any easier to solve than the unrestricted form of the problem. If the restricted form of the problem is easier to solve than its unrestricted cousin, we should expect to find the resources for a solution in those features that are unique to that problem. What is distinctive about the restricted form of the generality problem is that it applies only to cases of reflective reasoning – cases in which the *reasoner herself* identifies a BFPT as highly reliable and sets for herself the policy of using *that very BFPT*.

First Reply

Richard, our resident internalist, might attempt to solve the restricted form of the generality problem as follows: “When a person makes the explicit decision to reason in accordance with some BFPT, then it is the reliability of *that* BFPT that is relevant to determining the justificatory status of the belief it produces. This proposal has the right result in Heart Attack and Reliable Testimony. Sandra decided to use the heart attack model, and it is *that* BFPT (using the model) whose reliability is relevant to determining the justificatory status of her resulting belief. You decided to defer to *S* with respect to U.S. history, and it is *that* BFPT (deferring to *S* about U.S. history) whose reliability is relevant to determining the justificatory status of your Lincoln belief.”

On reflection, however, this reply gets Gerrymandered clearly wrong. Paula attempts to justify her absurdly optimistic beliefs by deciding to reason in accordance with a highly reliable BFPT. But despite appealing to this reliable BFPT, her beliefs are not justified. Richard's first proposed solution makes the justificatory status of a belief depend on how the believer decides to categorize the process-token that led to that belief. And that is a losing proposition because the generality problem teaches that any belief-forming process-token can be categorized as an instance of a reliable BFPT or as an instance of an unreliable BFPT. And so, in principle, a believer will always be able to cherry pick a BFPT that will make the justificatory status of her resulting belief be whatever she wants it to be.

Second Reply

Richard's second attempt to solve the restricted form of the generality problem runs as follows: "Sandra's target belief is justified because she knows that the heart attack BFPT is reliable – it usually produces true beliefs. And she knows that her belief about Sam was produced by a token process that is an instance of that BFPT. So her target belief is justified on simple inductive grounds: 'B was produced by a process that is an instance of P; most instances of P produce true beliefs; therefore, B is probably true.' That's not to say that Sandra necessarily engaged in this kind of explicit inductive inference. Rather, the point is that these inductive grounds are sufficient for Sandra's second-order beliefs to confer justification on the target belief."⁸

Once again, this reply fails because Richard's attempt to explain why the target belief in Heart Attack is justified can be applied, *mutatis mutandis*, to the target belief in Gerrymandered. Paula's belief can be supported on precisely analogous inductive grounds. But Paula's absurdly optimistic beliefs about her future are clearly unjustified, while Sandra's belief is clearly

⁸ John Pollock suggested this reply in conversation.

justified. Richard needs a *principled* way to explain why one reliable BFPT (the heart attack model) is relevant to the target's justificatory status while the other one (the gerrymandered BFPT) is not. Without a solution that is principled, pessimists (following Feldman) will accuse Richard of cherry picking – of first deciding whether the belief is justified and then identifying the BFPT that will get him the right answer.

Third Reply

Richard's final two replies attempt to elude the generality problem by insisting that in cases of reflective justification, second-order beliefs are epistemically impotent – they play no role in *justifying* target beliefs. As a result, the generality problem does not arise. The trick to pulling off this solution is to explain, given that the second-order beliefs do not justify the target belief, what *does* justify the target belief. Richard's third reply runs as follows: "The generality problem does not arise for internalism because in all cases of reflective justification, the reasoner's second-order beliefs are epistemically impotent. For example, Sandra's second-order beliefs (that the target belief was produced by the heart attack model and that the heart attack model is highly reliable) are epistemically impotent with respect to her target belief (that Sam is low risk). So what *does* justify Sandra's target belief? Her first-order beliefs about Sam (his age, his systolic blood pressure, etc.). Nothing else. The heart attack model was simply a useful tool that helped Sandra *identify* the belief justified by her first-order beliefs. Since Sandra's second-order beliefs about the BFPT play no role in justifying her target belief, I do not need a principled way to identify that BFPT as relevant to the justificatory status of Sandra's target belief. And so I don't need a solution to the generality problem. The same lesson holds in all cases of reflective justification. And so the generality problem never arises for internalism."

The problem is that while the reasoner's first-order beliefs will sometimes justify the target belief in cases of reflective justification, this will not always be true. This is particularly clear in Reliable Testimony. Your beliefs about Lincoln might be so vague or uninformed that they don't justify anything at all about his place in the sequence of U.S. Presidents. What justifies your Lincoln belief – or more precisely, what makes your Lincoln belief more justified than a run-of-the-mill testimonial belief – is your second-order knowledge that relying on S about U.S. history is a very reliable way to form beliefs (i.e., it is a very reliable BFPT). In Heart Attack, Sandra might come to believe that Sam is low risk even though she recognizes that her first-order beliefs about Sam (unaided by the heart attack model) do not support that belief. The entire *point* of the heart attack model is to relieve doctors of an onerous task. So Sandra might reckon that because she is so reliant on the model, she has not systematically searched for evidence relevant to Sam's risk status; and as a result, her first-order beliefs about Sam might be unrepresentative or incomplete in important ways. So the wise thing to do is to accept the model-based belief *because the model is very reliable, and in fact more reliable than she is – particularly given her failure to fully and systematically consider the available first-order evidence*. Sometimes, the second-order beliefs really are essential to the justification of the target belief.

At this point, I should mention a desperate maneuver for trying to save this third reply. Richard might bite the bullet by insisting that in cases of reflective justification, when a reasoner's first-order beliefs, by themselves, don't justify the target belief, the target belief is not justified. So in Heart Attack, if Sandra's belief that Sam is low risk is not justified by her first-order beliefs by themselves, then that belief is simply not justified. And in Reliable Testimony, if your belief that Lincoln was the 16th president of the U.S. is not justified by your first-order

beliefs by themselves, then that belief is not justified (or no more justified than a run-of-the-mill testimonial belief). But this is surely a mistake. When we're on the stretcher gasping for breath, we don't really believe that doctors are epistemically justified in ignoring the prediction models they know are most reliable. Or consider the models that predict violent recidivism more reliably than parole boards (Quinsey et al. 1998). We don't really want a theory of justification that would have us ignore them. There are models that predict more reliably than doctors whether a baby will die in the crib (Golding et al. 1985). We don't really believe that parents and doctors would be epistemically justified in ignoring the judgments of these models. At some point, epistemology leaves the armchair. There is no reason to make it a monster.

Fourth Reply

Richard's fourth reply, like the third, contends that second-order beliefs are epistemically impotent, and it runs as follows: "The internalist appeals to primitive epistemic principles in explaining why certain beliefs are justified. For example, if it seems to me as if there is a flower in front of me (and the background conditions are normal), then I am justified in believing that there is a flower in front of me. This is a *primitive* epistemic principle (or, perhaps more plausibly, it follows from primitive epistemic principles). As such, it does not explain the justification in terms of the alleged *reliability* of believing that there is a flower in front of me when it seems to me that there is a flower in front of me. Similarly, it is a primitive epistemic principle (or, perhaps more plausibly, it follows from primitive epistemic principles) that beliefs arrived at by following certain BFPTs (e.g., the heart attack model) are justified. However – and this is important – the primitive principle does not explain the justification of the target belief in terms of the *reliability* of the model or in terms of the reasoner's second-order beliefs. Otherwise, the principle would simply raise the specter of the generality problem once again:

Why is the reliability of *that* BFPT relevant to the justificatory status of the target belief but not the reliability of the indefinitely many other possible BFPTs?”⁹

To understand Richard’s fourth reply, it will be useful to reframe the generality problem in somewhat broader terms than is usual. The generality problem is usually described as a challenge for any attempt to justify a belief in terms of the *reliability* of the process that produced it: “A solution [to the generality problem] identifies the type whose *reliability* determines whether a process token yields justification” (Conee and Feldman 1998, 3, emphasis added). Let’s distinguish *relevant* BFPTs, BFPTs that are relevant to the epistemic evaluation of beliefs (like the decision tree in Heart Attack) from *non-relevant* BFPTs, BFPTs that are not relevant to the epistemic evaluation of beliefs (like the gerrymandered BFPT in Gerrymandered). The notions of relevant and irrelevant BFPTs are not defined in reliabilist terms. This is because both the reliabilist and the internalist need a principled way to distinguish relevant from non-relevant BFPTs. Keep in mind, Richard’s fourth reply explicitly assumes that Richard can identify the heart attack model as the BFPT that is *relevant* to evaluating Sandra’s belief that Sam is low risk. Richard disagrees with the reliabilist about what property of the heart attack model, and of BFPTs in general, is crucial to determining the justificatory status of Sandra’s target belief. So we can divide up the generality problem into two sub-problems:

1. Identify Relevant BFPTs: Provide a principled method that identifies relevant BFPTs. In order to be principled, (a) this method will be universal, in that it will identify a relevant BFPT whenever a BFPT is essential to the justification of a belief, and (b) it will not involve “*ad hoc* case-by-case selections of relevant types that match our intuitions” (Conee and Feldman 1998, 4).

⁹ The referee suggested this reply.

2. Identify the property, J, of relevant BFPTs that is crucial to determining the justificatory status of beliefs produced by those BFPTs.¹⁰

So here's the state of the debate: To solve the generality problem, the reliabilist only needs a solution to problem 1; the reliabilist has a ready answer to problem 2 – it is the reliability of the BFPT that is crucial to determining the justificatory status of B. Richard's fourth reply commits him to solving both problem 1 and problem 2; he adverts to an unarticulated principle, let's call it UP, that he envisions will do this.

With this broader characterization of the generality problem in mind, Richard's fourth reply suffers from three problems. First, it has all the advantages of theft over honest toil (to borrow a phrase from Russell). For decades now, internalists have argued that without a solution to the generality problem, reliabilism is in serious trouble. It takes some chutzpah for the internalist to dismiss essentially the same problem armed with nothing but the promise of an as-of-yet unarticulated epistemic principle. Second, the reliability problem facing the internalist appears to be at least as difficult as the reliability problem facing the reliabilist. Both have to come up with a solution to problem 1, but only the internalist still needs a solution to problem 2. And it's going to be tough for the internalist to solve problem 2. We know that for Richard, J, the property of relevant BFPTs that is crucial to determining the justificatory status of beliefs produced by those BFPTs, cannot be *reliably produces true beliefs*. That would put Richard back at square one. J also cannot be *reliably produces beliefs that are best supported by (or best fit) the subject's first-order beliefs*. As we have already seen in the discussion of Richard's third

¹⁰ Nothing of substance turns on this broader characterization of the generality problem. Those who want to insist that the generality problem applies by definition only when J is the reliability of the relevant BFPTs can call the broader problem *the generality problem**. Given that way of describing things, my thesis is that all plausible theories of justification suffer from the generality problem* and this problem is at least as knotty as the generality problem.

reply, even if your first-order beliefs about U.S. history are deeply confused, you are justified in relying on S's views about U.S. history. But the beliefs that result from this BFPT very well might not be reliably supported by your confused beliefs about U.S. history. This point – that the problem facing the internalist is at least as difficult as the problem facing the reliabilist – is crucial for my purposes. That's because my thesis is *not* that the generality problem *defeats* internalism. I am not committed to the claim that the internalist cannot come up with some epistemic principle that will solve the generality problem. I am arguing simply that there is no reason to think that the internalist's generality problem is easier to solve than the reliabilist's generality problem.

The third problem with this reply is that if Richard succeeds in coming up with a solution to problems 1 and 2, *the reliabilist can simply co-opt Richard's solution to problem 1*. In other words, any principle that would solve the generality problem for the internalist would solve the problem for the reliabilist as well. To see this, let's return to Heart Attack. UP (by hypothesis) will identify the heart attack model as the relevant BFPT for Sandra's target belief. For any case of reflective justification, the reliabilist can co-opt UP to come to this same conclusion. But what about other beliefs – beliefs not justified reflectively? The reliabilist can co-opt UP for these beliefs as well. Let's consider a case that is analogous to Heart Attack but that does not involve reflective justification. Andres, as a psychological matter of fact, is naturally disposed to employ the heart attack model whenever faced with a heart attack patient. His default reasoning mechanism instantiates the model, but he does not know how his default reasoning mechanism works or how reliable it is. Suppose that Andres acquires the same first-order beliefs about Sam that Sandra does, and as a result of the operation of his default reasoning mechanism Andres comes to believe that Sam is low risk, as does Sandra. Andres's case raises an instance of the

standard, unrestricted form of the generality problem. It is a problem only for the reliabilist, not for the internalist: Why is the decision tree the relevant BFPT in Andres's case? I am not proposing an answer to this question; I am simply assuming that Richard has a principled answer to this question *in Sandra's case*. My contention is that the reliabilist can apply Richard's answer from Sandra's case directly to Andres's case. *That's because the only differences between Sandra and Andres are differences that Richard's unarticulated principle must ignore*. Consider the facts that distinguish Sandra and Andres:

- a. Sandra intentionally chose to use the heart attack model whereas Andres is naturally disposed to use it. As a principled method of identifying relevant BFPTs, UP cannot employ Sandra's intentionally choosing to use the heart attack model to explain why that model is the relevant BFPT. This would mistakenly identify Paula's gerrymandered BFPT as relevant to the epistemic evaluation of Paula's absurdly optimistic beliefs about her future. (This point was made in response to Richard's first reply.)
- b. Sandra has (second-order) knowledge about the workings of the heart attack model and its reliability whereas Andres does not. As a principled method of identifying relevant BFPTs, UP cannot employ Sandra's second-order knowledge to explain why the heart attack model is the relevant BFPT. This would mistakenly identify the gerrymandered BFPT as relevant to the epistemic evaluation of Paula's unjustified beliefs. Furthermore, this explanation seems inconsistent with Richard's fourth reply, which explicit says that Sandra's second-order beliefs are irrelevant to determining the epistemic status of the target belief.
- c. Sandra's second-order beliefs appear to provide inductive support to the target belief. UP cannot appeal to this fact about Sandra in a principled explanation for why the heart

attack model is the relevant BFPT. This explanation, again, gives the wrong answer in Gerrymandered, and it seems inconsistent with Richard's fourth reply, which is explicitly committed to the epistemic impotence of Sandra's second-order beliefs.

So take any instance of *reflective* justification in which S1 justifies belief B by appeal to belief-forming process-type T. UP will offer a principled explanation for why T is relevant to the epistemic evaluation of S1's belief B. I have argued that the reliabilist will always be able to co-opt that principle to explain an analogous instance of non-reflective justification – where S2 acquires B via T but knows nothing about the operation of T and so does not explicitly justify B by appeal to T. The reliabilist will be able to use UP to give a principled explanation for why T is relevant to the epistemic evaluation of S2's belief B. Since any belief can, in principle, be justified by reflection, UP solves problem 1 for the reliabilist. And since the reliabilist needs only to solve problem 1 to solve the generality problem, UP solves the generality problem for the reliabilist.

Recall from section 1 what makes the generality problem hard for the reliabilist – she has to come up with a *principled* way to pick out the *single* BFPT that is relevant to the justificatory status of the target belief out of *indefinitely* many non-relevant BFPTs. Richard's fourth reply envisions a solution to this exact problem. Once the reliabilist has a solution to problem 1, the generality problem is solved. Richard's fourth reply, therefore, expresses a deep optimism about the prospects for solving the generality problem. But as we have seen, many philosophers reject this optimism. They think the prospects for finding a principle that will solve the generality problem for the reliabilist are “worse than bleak” (Conee and Feldman 1998, 5). But if the pessimists are right, reliabilism is not the only victim. Internalist theories of justification, including the evidentialism of Feldman and Conee (1985), “look hopeless” as well (Conee and

Feldman 1998, 24). With respect to the generality problem, all plausible theories of justification sink or swim together.

5. Conclusion

If the generality problem has no solution, then no plausible theory of justification – no theory that accounts for reflective justification – survives. Given this conditional, we can either tollens or ponens: It is either time to give up trying to develop theories of epistemic justification or it is time to stop being pessimistic about the prospects for solving the generality problem.

Philosophers have proposed quite a few solutions to the generality problem (e.g., Goldman 1986, Schmitt 1992, Wallis 1994, Alston 1995, Heller 1995, Wunderlich 2003, Comesaña 2006). It is fair to say that none of them have met with widespread support (e.g., Conee and Feldman 1998; Feldman and Conee 2002). At the risk of sounding alarmist, maybe it is time to think about the implications for epistemology if the pessimists about the generality problem are right.

Acknowledgements

For many excellent questions on earlier versions of this paper, I would like to thank audiences at the University of Pennsylvania, Stockholm University, the University of Gothenburg, and at the conference on The Epistemology of Natural and Artificial Systems at California State University, Long Beach. I would also like to thank Stewart Cohen, Hilary Kornblith, John Pollock, J.D. Trout, Charles Wallis, an anonymous referee, and especially Mark Wunderlich for very helpful comments on earlier drafts of this paper. I am also grateful to the National Science Foundation for grant SES#0354536 that supported this research in its early stages.

References

- Adler, J. and M. Levin. 2002. Is the Generality Problem too General? *Philosophy and Phenomenological Research* 65, 1: 87-97.
- Alston, W. 1995. How to Think about Reliability. *Philosophical Topics* 23, 1: 1-29.
- Breiman, L., J. Friedman, R. Olshen, C. Stone. 1984. *Classification And Regression Trees*.
Monterey, CA: Wadsworth.
- Comesaña, J. 2006. A Well-Founded Solution to the Generality Problem. *Philosophical Studies* 129: 27-47.
- Conee, E. and R. Feldman. 1998. The Generality Problem for Reliabilism. *Philosophical Studies* 89: 1-29.
- Feldman, R. 1985. Reliability and Justification. *The Monist* 68: 159-174.
- Feldman, R. and E. Conee. 1985. Evidentialism. *Philosophical Studies* 48: 15-34.
- Feldman, R. and E. Conee. 2002. Typing Problems. *Philosophy and Phenomenological Research* 65, 1: 98-105.
- Golding, J., S. Limerick, and A. MacFarlane. 1985. *Sudden Infant Death*. Somerset: Open Books
- Goldman, A. 1979. What is Justified Belief? In *Justification and Knowledge*, ed. G. Pappas.
Dordrecht: Reidel.
- Goldman, A. 1986. *Epistemology and Cognition*. Cambridge, MA: MIT Press.
- Heller, M. 1995. The Simple Solution to the Generality Problem. *Nous* 29, 4: 501-515.
- Lackey, J. 2006. Introduction. In *The Epistemology of Testimony*, eds. J. Lackey and E. Sosa.
Oxford: Oxford University Press.
- Pollock, J. and J. Cruz. 1999. *Contemporary Theories of Knowledge, 2nd ed.* Lanham, MD:
Rowman and Littlefield.

- Quinsey, V. L., G.T. Harris, M.E. Rice, and C.A. Cormier. 1998. *Violent Offenders: Appraising and Managing Risk*. Washington, D.C.: American Psychological Association.
- Schmitt, F. 1992. *Knowledge and Belief*. London: Routledge.
- Sosa, E. 1985. Knowledge and Intellectual Virtue. *The Monist* 68: 224-245.
- Shafer-Landau, R. 2003. *Moral Realism: A Defence*. Oxford: Oxford University Press.
- Wallis, C. 1994. Truth-Ratios, Process, Task, and Knowledge. *Synthese* 98: 243-269.
- Wunderlich, M. 2003. Vector Reliability: A New Approach to Epistemic Justification. *Synthese* 136: 237-262.