Philosophical Studies (2006) 129:27–47
DOI 10.1007/s11098-005-3020-z

JUAN COMESAÑA

A WELL-FOUNDED SOLUTION TO THE GENERALITY PROBLEM

1. INTRODUCTION

Consider reliabilism about epistemic justification:

Reliabilism: a belief is epistemically justified if and only if it was produced by a reliable belief-forming process.

It is a matter of controversy exactly what it is for a belief-forming process to be reliable. But the general idea is that a belief-forming process is reliable just in case it issues in true beliefs most of the times it is used – or it *would* issue in true beliefs most of the time if it *were* used under appropriate circumstances.[1]

Earl Conee and Richard Feldman have forcefully presented a problem for reliabilism, "the generality problem."[2] The generality problem arises once we realize that the notion of reliability applies straightforwardly only to *types* of process – for only types of process are repeatable entities which can produce true or false beliefs in each of their instances. Moreover, any token process will be an instance of *indefinitely many* types of process – for instance, the token process that resulted in my belief that there is a computer in front of me is an instance of all of the following types of process: a process initiated in a visual image, a process that starts with an experience as of a computer in front of me and ends with the belief that there is a computer in front of me, a process that takes place on Monday, a process that takes place while I am wearing black socks, etc. Which of these types must be reliable for my belief to be justified, according to reliabilism?

That question, generalized to every case of belief-formation, is the generality problem for reliabilism.

In this paper I propose a solution to the generality problem. The solution makes use of the *basing relation*, and so, given that it isn't clear how to characterize that relation, it might be thought to replace one problem with another. I argue that, however difficult it is to characterize the basing relation, every adequate epistemological theory must make use of it implicitly or explicitly. Therefore, it is perfectly legitimate to appeal to the basing relation in solving a problem for an epistemological theory.

Before presenting my own proposal, I shall first argue against the idea that we can avoid the generality problem altogether simply by focusing on token processes.

2. A QUICK FIX?

In explaining the generality problem, I said that the notion of reliability applies *straightforwardly* only to types of process. But maybe there is a less straightforward fashion in which the notion of reliability applies to process-tokens as well. One might wonder, then, whether reliabilists can escape the generality problem by saying that the relevant reliability is that of the *token* process that issued in the belief. Given that there is only *one* such token process, no generality problem arises.[3] I don't think that token reliabilism can solve the generality problem, but not for the reason that might immediately spring to mind.

The version of reliabilism that we have to consider is the following:

Token reliabilism: a belief is epistemically justified if and only if it was produced by a reliable *token* process.

One might think that token reliabilism is a non-starter because there is no sense in asking what truth-value *most of* the outputs of a token process will have or would have. A process-type is a repeatable entity, and thus it makes sense to ask *what proportion* of the token processes of a given type has

true outputs, and what the result would be if a process of the given type were instantiated. But a *token* process is an individual, a dated entity. As William Alston has put it, "a particular process that takes place at a particular precise time is not the sort of thing that does or does not enjoy a favorable ratio of true beliefs among its products."[4]

But this objection to token reliabilism is misguided. To see why, consider particular persons or objects, such as yourself or the chair you are sitting on. Both are unrepeatable. But *of course* it makes sense to ask what would have happened to you had you been in different circumstances, or what would have happened to that chair had it been in the other room. And just as we can meaningfully ask what would have happened to you had you (*you*, not merely someone very much like you) pursued a career in molecular biology, we can also meaningfully ask what truth value my belief would have had if the token process that produced it (*that very same process*, not merely a different process of the same type) had operated in different circumstances. And that is all we need to make sense of token reliabilism – for, once we have token processes that could have had different properties (token processes that, so to speak, inhabit different possible worlds), we do have "the sort of thing that can enjoy a favorable ratio of true beliefs among its products." The ratio will not be measured only over the process's *actual* outcomes, of course (it has only one of those), but over (some of) its possible outcomes as well.

My defense of token reliabilism from the objection that it doesn't make sense to talk of the reliability of a token appeals to the idea that a token process can inhabit different possible worlds – that it could have operated under different conditions. Token reliabilism will not make sense if one has a view of processes according to which they are extremely "modally fragile": the view that a token process can only have the exact same properties as it actually has, and any other process with even slightly different properties would have been a different process. But this view about processes is *prima facie* implausible. Intuitively, *that fire* could have burned that house, although it didn't, and *that*

election could have had a different result. That is, it is intuitively right to say of many processes that they (those very same processes) could have had different properties from the ones they actually have.

It does make sense, then, to talk of the reliability of a token process. But that is not sufficient for token reliabilism to be adequate. For in order to determine whether a given token process is reliable, we have to select a subset of those worlds that contain the token process and see what the ratio of true to false beliefs is in those worlds.[5] But this just is the generality problem under a different guise.[6] If we have a selection of the possible worlds where I use the process that contains all and only those worlds that are relevant to determining the reliability of the process, then we have a solution to the generality problem: the type that is relevant in determining the reliability of the process is the type determined by that subset of possible worlds.[7] Appealing to token processes doesn't provide us with a quick fix to the generality problem.[8]

3. THE WELL-FOUNDEDNESS PROBLEM

In this section I argue that there is a problem, the problem of well-foundedness, that any plausible epistemological theory must face. In the next section I will argue that a solution to the well-foundedness problem gives us the elements to solve the generality problem.

In order to understand the well-foundedness problem, it will be useful to introduce a distinction that is customary to make in discussions of epistemic justification. Theories of epistemic justification can be divided into two general categories, according to whether the central notion of the theory is that of a proposition being justified for a subject, whether the subject actually believes it or not, or that of an episode of believing being justified for a subject. Let's call theories of the first kind, theories of *propositional* justification, and theories of the second kind, theories of *doxastic* justification.

Reliabilism is clearly a theory of doxastic justification, for only episodes of believing (and not propositions) are produced

by processes. More generally, externalist epistemic theories tend to be theories of doxastic justification.[9] By contrast, internalist theories are best conceived as theories of propositional justification.[10] As we will see, the well-foundedness problem arises most clearly in connection with theories of propositional justification, because theories of doxastic justification incorporate the idea of well-foundedness into the notion of justification itself.[11] Indeed, the generality problem arises for reliabilism precisely because it is a theory of doxastic justification – for the problem with the characterization of well-foundedness is intuitively the same as the generality problem, and so any theory of justification that incorporates the notion of well-foundedness into the characterization of justification will have to deal with the generality problem. Theories of propositional justification avoid the need to face the generality problem head-on only because they do not incorporate the notion of well-foundedness into the notion of justification itself. But the problem cannot be avoided for long, for a theory of propositional justification without a theory of well-foundedness is incomplete as an epistemological theory.

Let us see, then, how the problem of well-foundedness arises for one specific theory of propositional justification, Conee and Feldman's own "evidentialism." Evidentialism is the theory that for a proposition to be epistemically justified for a subject is for it to fit the evidence that the subject has. More specifically:

Evidentialism: Doxastic attitude D toward proposition p is epistemically justified for S at t if and only if having D toward p fits the evidence S has at t.[12]

It is obvious that evidentialists must clarify what it is for a belief to "fit" the evidence and what it is for a subject to "have" some evidence,[13] but let us leave that aside. With respect to what counts as evidence, Feldman and Conee think that sensory states (such as having a visual experience as of seeing blue) as well as beliefs count as evidence.

Evidentialism as stated is open to the following kind of counterexample: I believe that there is a computer in front of

me, and my sensory state is that of seeming to perceive a computer in front of me, but I have the belief because I see something gray in front of me, and whenever I see something gray in front of me I believe that there is a computer in front of me. Intuitively, my belief is not justified, although it does fit my evidence.

We should not underestimate the importance of this kind of counterexample to evidentialism. Without an answer to the problem posed by this kind of case, evidentialism is simply bankrupt as an epistemological theory: evidentialism clearly implies that my belief is justified, and my belief is clearly not justified. Given this, Conee and Feldman's reaction to this kind of counterexample is surprising. Instead of modifying the theory in the light of cases like this, they choose to say that my belief about the computer *is*, after all, justified, although it is not "well-founded." Before saying more about this notion of well-foundedness, it is worth noting one thing about this move.

We should grant, I think, that the proposition that there is a computer in front of me has, in the case as described, something good going for it, epistemically speaking. But the epistemic goodness enjoyed by this proposition is only an *impersonal* kind of epistemic goodness – the kind of epistemic goodness that old-fashioned philosophers of science were after, where the central notion is that of a scientific theory's being justified, regardless of the reasons that any particular scientist has for believing it (and, indeed, regardless of whether any scientist believes it). And it is all well and good to be concerned with impersonal justification, of course, but it is hardly undeniable that a central concern of epistemology historically has been (and should continue to be) *personal* justification, where what matters are the reasons that a subject has for holding a certain belief. Because of this, any theory of propositional justification that doesn't incorporate as well a theory of well foundedness will simply be incomplete as an epistemological theory.

Let us now look more closely into the notion of well-foundedness. The intuitive idea is that for a belief to be well-founded

it has to not only fit the subject's evidence, but also *be based on* that evidence. Conee and Feldman define the notion of well-foundedness as follows:

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

(i) having *D* toward *p* is justified for *S* at *t*;
and
(ii) *S* has *D* toward *p* on the basis of some body of evidence *e*, such that

(a) *S* has *e* as evidence at *t*;

(b) having *D* toward *p* fits *e*; and

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

It is important now to note that there are good reasons to think that every adequate epistemological theory, and not just evidentialism, is going to have to appeal to the basing relation. Every adequate epistemological theory had better have the consequence that my belief that there is a computer in front of me in the case described above is not justified – or, if the term "justified" is to be reserved for a property that my belief does have in that case, then the theory had better have the consequence that there is some positive epistemic property that my belief lacks. And it seems that the most natural way to say what it is that my belief lacks in the case as described is to say that it is not adequately based on the factors that would make it justified. And this is, in effect, to appeal to the basing relation. This appeal might be implicit, as I think it will be in most instances of adequate doxastic theories, but it will still be there if the theory has the resources to handle the problematic cases. It follows that we have good reasons to believe that any adequate epistemological theory needs to appeal, either implicitly or explicitly, to the notion of a belief's being based on certain evidence. In the next section I will argue that that notion is all we need to solve the generality problem. If this is so, then any adequate epistemological theory is going to have the resources to solve the generality problem.[14]

4. A SOLUTION TO THE GENERALITY PROBLEM

Many philosophers have tried to solve the generality problem by saying, roughly, that the relevant process-type to assess for reliability must be chosen taking into account the actual production of the belief. Thus, William Alston holds that there is a fact of the matter about what psychological features in the mental life of the subject had an impact in the formation of the belief in question, and what kind of impact they had. There will always be, then, a process-type-schema, *having been produced by such and such mental antecedents in such and such a way*, that we can fill out in each case in order to get a specific process-type to assess for reliability. In Alston's words:

> [E]very belief formation involves the activation of a certain psychologically realized *function*... The function involved will determine both what features of the input have a bearing on the belief output and what bearing they have.[15]

Conee and Feldman have replied to Alston, and Jonathan Adler and Michael Levin have come to Alston's defense – and Conee and Feldman have in turn replied to Adler and Levin.[16] In what follows I will make some comments regarding this discussion of Alston's proposal, and I will end by arguing that Conee and Feldman cannot deny that *something like* Alston's proposal has to work.

In their reply to Alston, Conee and Feldman granted that Alston's *psychological realism* – the claim that there is a fact of the matter about what features of the input have a bearing on the output belief – does restrict the number of types that, for all we know, are relevant to assessing the reliability of the token process that generated the belief. They discuss a case where Smith faces a tree and forms the belief that it is an apple tree. "Thus," they say, "in our example about Smith and the apple tree, Smith might form her belief on the basis of noticing certain features of leaf shape," and this "show[s] that the relevant type in the original case must be one that corresponds to a function having as input/output pair the leaf-

shape features to which Smith responds and the belief that she forms." However, they insisted that, even with this restriction, there is still more than one type that the process falls under:

> There is a very narrow function that goes from just the leaf shape that Smith notices as input to just the output of Smith's particular belief that a maple tree is nearby. There is another function, one that maps a variety of fairly similar inputs, including the particular shape that Smith noticed, onto some belief or other to the effect that there is a maple tree nearby, including the belief Smith forms. There is a broader function, one that maps a variety of somewhat similar inputs, all involving visual shapes, onto either the belief that there is a maple tree nearby or the belief that there is an oak tree nearby or the belief that there is an elm tree nearby, etc. There are still broader types that include the original pair, and add new inputs involving various other sensory cues. In many cases, all these functional causal relations, and many others as well, would be actually operative in forming Smith's belief. Smith's disposition to form the particular belief that she did on the basis of the particular shape that she saw is part of these broader classifying dispositions. The one event of belief-formation manifests them all. Thus, in this and other typical cases, there are a multitude of actually operative psychological types.[17]

Adler and Levine reply, on behalf of Alston, that Conee and Feldman are not really describing different functions, but rather the same function in increasing level of detail. More precisely, Adler and Levine think that some of Conee and Feldman's descriptions are really different descriptions of the same function, and others do indeed describe different functions, but these different functions are *not* all of them "actually operative in forming Smith's belief." For example, suppose that the subject believes that it is a maple tree because its leaves have seven points. In that case, whereas Conee and Feldman would say that the token process that produces the subject's belief is an instance of many different function-types, e.g., the function that consists *just* in the ordered pair $<7, \text{maple}>$, and the function $\{<7, \text{maple}>, <9, \text{oak}>\}$, Adler and Levine think that the function partially described by the ordered-pair $<7, \text{maple}>$ is the same as the function that is more fully (although maybe still partially) described by the ordered-pairs $<7, \text{maple}>$ and

<9, oak$>$. There is only one function, one that takes number of points in a leaf to types of trees. On the other hand, the function partially described by $<<7$, brown$>$, maple$>$ (which corresponds to the process that takes as inputs the fact that the leaves have seven points and the bark is brown to the belief that it is a maple tree) *is* different from the previous one, but "no instance of the process of forming beliefs from point number *alone* also instances the process of forming beliefs from point number and further factors such as bark color."[18]

Predictably, Conee and Feldman disagree.[19] They think that the function partially described by <7, maple$>$ *need not be* the same as the function partially described by <9, oak$>$. I don't want to enter too much into this dispute, because, as we will see shortly, I think that it can be bypassed, but let me briefly say why I think that all of Adler, Levine, Conee and Feldman are wrong. It seems to me that what Adler and Levine say about the last function is also what they should have said about the previous one. That is, I think that Conee and Feldman are right in saying that the function partially described by <7, maple$>$ *need not be* the same as the function described partially described by <9, oak$>$. But I think that they are wrong in thinking that, in normal cases, a token process will instantiate both functions. That the functions partially described by those two ordered pairs need not be the same is obvious: take the function $\{<6$, maple$>$, <7, maple$>\}$ and the function $\{<8$, oak$>$, <9, oak$>\}$, for example. What's more, even if we restrict the domain to functions that are descriptions of possible processes of belief formation, it seems entirely plausible to suppose that a process can instantiate one function without instantiating the other – just think of a subject quite knowledgeable about maples but oblivious to the existence of oaks. But, for exactly this reason, it seems to me that Conee and Feldman are wrong in saying that "[i]n many cases, all these functional causal relations, and many others as well, would be actually operative in forming Smith's belief."

I do not presume for a second that this is the end of the dialectic. But, as promised, I will now argue that we do not even need to enter into the dialectic in the first place, for Conee and Feldman, and every other philosopher who wants to have an adequate epistemological theory, are *committed* to the existence of a solution to the generality problem.

Remember that Conee and Feldman are committed to the existence of a *basing relation* that obtains between the evidence that a subject has and his belief. When the evidence justifies the belief and the subject bases his belief on this evidence, the belief is well-founded – and this is a crucial ingredient of Conee and Feldman's reply to the cases that are problematic for evidentialism. It could of course happen that a belief is based on evidence that is not good enough to justify the proposition in question – never mind, in that case too the evidence is related to the act of believing by the basing relation.

The key to solving the generality problem is an appeal to this basing relation that figures in the account of well-foundedness. Given that there will always be some evidence that the belief is based on, the process that generates the belief will always instantiate a case of the type-schema *producing a belief that p based on evidence E*. This is only a type-*schema* because it contains the variables *p* and *E*, but in each particular case of belief formation we will have a particular instance of this schema: *producing a belief that the tree over there is an elm based on the fact that it looks as if the leaf has seven points*; or *producing a belief that it is cloudy based on the belief that it is raining and the belief that if it is raining then it is cloudy*; etc. The proposal, then, is that those are the types that we have to assess for reliability when applying reliabilism as a theory of epistemic justification.[20] So, if the belief that the tree over there is a maple was based on the fact that it looks as if its leaves have seven points, then, according to reliabilism, the belief is justified if and only if the process *producing a belief that the tree over there is an elm based on the fact that it looks as if the leaf has seven points* is a reliable

process of belief formation, and the same goes, *mutatis mutandis*, for other processes and other beliefs.

More explicitly, the proposed solution, then, is the following:

Well-Founded Reliabilism: A belief that *p* by *S* is epistemically justified if and only if:

(i) *S* has evidence *E*;

(ii) the belief that *p* by *S* is based on *E*; and

(iii) the type *producing a belief that p based on evidence E* is a reliable type.

This solution should be accepted not only by Conee and Feldman, but by anyone who thinks that an epistemological theory is incomplete without an (explicit or implicit) appeal to the basing relation. And, as I argued above, there are strong reasons to believe that any epistemological theory *is* incomplete without an appeal to the basing relation.

5. THE ADEQUACY OF THE WELL-FOUNDED SOLUTION

Conee and Feldman say that an adequate solution to the generality problem must satisfy three conditions: (i) it should be principled, rather than an ad-hoc, case-by-case selection of a relevant type; (ii) it must make defensible epistemic classifications (counting as justified all and only those beliefs that we would count as justified pre-theoretically); and (iii) it must be true to the spirit of reliabilism. I think that these are fair adequacy conditions on a solution to the generality problem, and I will next argue that the well-founded solution does satisfy them.

It is quite clear that the well-founded solution satisfies condition (i). For every belief, it provides a *general* recipe to recover the relevant type to be assessed for reliability, rather than relying on characteristics that are specific to the case at hand.

I also think that the well-founded solution satisfies condition (ii), at least on one interpretation of reliability. The most worrisome range of cases comes from a traditional objection to reliabilist theories of justification, an objection

that has been called "the new evil demon problem."[21] The problem is that a victim of an evil demon (or a brain in a vat, or...) who is otherwise a responsible epistemic agent (doesn't commit more fallacies than us, takes experience at face value, etc.) seems to be just as justified as we are, and yet her beliefs are not produced by reliable processes – or so the objection goes. This is not the place to develop in detail an answer to this worry, but I do think that it can be answered by requiring that the process be *actually* reliable.[22]

Does the well founded solution satisfy condition (iii) – that is, is it "true to the reliabilist spirit"? I'm not sure what it takes for a solution to be true to the reliabilist spirit, but the well-founded solution is *not* equivalent to one that Conee and Feldman criticize for just this reason. They say that

> one could develop a form of "reliabilism" that just restates an evidentialist theory of justification in a roundabout way. Pseudo-reliabilism of this sort holds that there are only two relevant types of belief-forming process. One type is "belief based on adequate evidence" and the other type is "belief based on inadequate evidence." Assuming that the first of these is reliable and the second is not, this version of reliabilism will get plausible results (or at least results that an evidentialist would find plausible). But this theory is only verbally a version of reliabilism. It mentions the process of belief-formation only in order to characterize the quality of the evidence for the belief. This is obviously incompatible with the spirit of process reliabilism.[23]

My proposed solution does *not* consist in saying that a belief is justified if and only if it is based on good evidence. My proposal, instead, suggests a way of characterizing what it is for a piece of evidence *to be* good evidence that a proposition p is true – it is for it to be the case that, if you base your belief that p on that evidence, then your belief would tend to be true. The notion of *evidence* itself is not, of course, the exclusive property of evidentialism. Indeed, we can think of a subject's evidence for a given belief as whatever psychological features that subject bases the belief on, and then characterizing what it is for the evidence to be *good* along the reliabilist line just proposed – this would be one natural way

of saying what it is for a given piece of evidence to "fit" a certain belief.[24]

There is another worry with the well-founded solution that I wish to discuss.[25] Suppose that both Ms. Accurate and Mr. Sloppy are in the following situation: they are facing a many-sided black figure against a white background under good lighting conditions (and while both are sober, etc.). The figure has 386 sides. Both Ms. Accurate and Mr. Sloppy have the capacity, to some degree or other, of telling how many sides a figure has just by briefly looking at it (without counting). For Mr. Sloppy, as for you and me, this capacity gets worse when the number of sides reaches five or six, and is just completely worthless when there are more than, say, eleven sides. But Ms. Accurate has the capacity to an extraordinary degree: she can tell how many sides a figure has just by briefly looking at it for figures with up to a thousand sides.

Now, Ms. Accurate, making use of this capacity, forms the belief that the figure in front of her has 386 sides. Mr. Sloppy can't really tell how many sides the figure has, but he nevertheless ventures the guess that it has 386 sides. Intuitively, Ms. Accurate is justified in her belief, but Mr. Sloppy is not justified in his belief. But the well-founded solution entails that both of their beliefs have the same justificatory status, because they are both produced by a process that falls under the type *producing a belief that the figure in front of you has 386 sides based on how it looks to you*. This type is either reliable enough or it isn't. If it is, then the well-founded solution entails that Mr. Sloppy is justified, and, if it isn't, then the well-founded solution entails that Ms. Accurate is not justified. In any case, the well-founded solution has a counterintuitive consequence.

Before replying to this objection, it is important to notice that, if it works, it applies just as much to evidentialism as it does to well-founded reliabilism – for the objection is, really, an objection to the definition of well-foundedness that we appealed to before. Both Mr. Sloppy and Ms. Accurate satisfy this definition with respect to the belief that the figure in

front of them has 386 sides, and yet, intuitively, only Ms. Accurate's belief is well-founded – not because Mr. Sloppy's belief is based on a different evidence (it isn't), but because it is based on that same evidence in a defective way.[26] So, the problem here is really a problem with the notion of well-foundedness itself, not with the specific use of it in well-founded reliabilism.

There are at least two ways in which a defender of the well-founded solution can reply to this objection. First, it can be denied that Ms. Accurate and Mr. Sloppy have the same evidence for their belief. Notice that the evidence is *not* the fact that the figure that they are facing has 386 sides, but rather their respective experiences as of a figure with a number of sides. And it can be plausibly denied that these experiences are the same for Ms. Accurate and Mr. Sloppy. After all, one natural way of explaining Ms. Accurate's superior discriminating abilities is by supposing that, when faced with a figure with 386 sides, her visual experience is much more detailed than Mr. Sloppy's experience when *he* is faced with the same figure. Ms. Accurate can tell just by sight how many sides the 386-sided figure has because the fact that it has 386 sides is clearly represented in her experience, whereas we (including Mr. Sloppy) cannot tell because our experience only represents the fact that the figure has a lot of sides. If this explanation of Ms. Accurate's superior abilities is adequate, then the comparison between her and Mr. Sloppy no longer represents a problem for well-founded reliabilism. For now the relevant process types that have to be evaluated for reliability will be different. In Ms. Accurate's case it will be the type *producing a belief that the figure in front of you has 386 sides based on the fact that it looks to you as if it has 386 sides*; whereas in Mr. Sloppy's case the relevant type will be *producing a belief that the figure in front of you has 386 sides based on the fact that it looks to you as if it has a lot of sides*. These two types clearly differ in reliability, and, moreover, in a way that dovetails nicely with our intuitions regarding the cases.

So the objector to well-founded reliabilism has to insist that, at least in some cases, someone can have a superior

discriminatory ability that *cannot* be explained in terms of the fact that the same things look differently (more clearly) to that subject. For the objection to work, the superior ability has to reside in the *way* of transitioning from the *same* evidence as the rest of us to an accurate belief. Now, the defender of well-founded reliabilism could at this point dig in his heels and insist that any superior discriminatory ability *has* to be explained in terms of a capacity to acquire better evidence. This is, I think, ultimately an empirical question.[27] But I think that, even if the world doesn't make the well-founded reliabilist the favor of making it the case that superior discriminatory abilities can always be explained in terms of a capacity of acquiring better evidence, there is still another answer available.

If Ms. Accurate's has the same experiential input as Mr. Sloppy, then her superior capacity can only be explained in terms of *how* she gets from this input to the belief. But, if that is the case, then we can just build this way of getting from the input to the belief into the description of the type. The relevant type-schema will then no longer be *producing a belief that p based on evidence E*, but rather *producing a belief that p based on evidence E by using method M*. Applied to our case, this would yield the processes *producing a belief that the figure in front of you has 386 sides based on the fact that it looks to you as if it has a lot of sides by using the method of discriminating exactly how many sides it appears to have* (for Ms. Accurate), and *producing a belief that the figure in front of you has 386 sides based on the fact that it looks to you as if it has a lot of sides by using the method of venturing a wild guess as to how many sides it has* (for Mr. Sloppy). Again, the reliability that is plausibly associated with these two methods coheres nicely with our intuitions about the case.

Someone could retort that this second answer to the objection makes the generality problem appear once again – for, how are we going to choose the right description of the method to use in the specification of the relevant type? After all, the *token* method that Ms. Accurate used in order to get from the experience to the belief falls under indefinitely many types:

not only *discriminating exactly how many sides it appears to have*, but also *discriminating how many sides it appears to have while wearing black socks*, and, indeed, *believing that it has 386 sides when it looks as if it has a lot of sides* (and many more). And the same goes, of course, for the method that Mr. Sloppy used. So it seems that we haven't advanced much in our attempt to solve the generality problem.

At this point it is important to remember that the problem is, fundamentally, a problem for the notion of well-foundedness itself. Ms. Accurate's belief is not well-founded in virtue of the fact that she transitioned from her experience to her belief by way of discriminating how many sides the figure appears to have while wearing black socks, but rather it is well-founded in virtue of the fact that she made the transition by way of discriminating how many sides the figure appears to have. If we are to have an adequate theory of well-foundedness, then we need to be able to discriminate between the processes that, intuitively, are relevant to assessing well-foundedness from those that are not. Once that work is done (if it can be done), the well-founded reliabilist can make use of it. And, if that work cannot be done, then we have much more serious problems than the generality problem.

6. CONCLUSION

Conee and Feldman end their paper on the generality problem by saying:

> That is the full variety of existing approaches to the generality problem. In the absence of a brand new idea about relevant types, the problem looks insoluble. Consequently, process reliability theories of justification and knowledge look hopeless.[28]

I don't think that the well-founded solution to the generality problem is a brand new idea about relevant types. But I do think that by making an explicit appeal to the basing relation, the well-founded solution highlights that the resources needed to solve the generality problem must already be

present in any adequate epistemological theory. Of course, it might turn out that there are insurmountable difficulties in giving an account of what it is for a belief to be *based on* a certain piece of evidence – but reliabilism would not be the only theory to suffer from this outcome. So Conee and Feldman's pessimism about the generality problem is at best unwarranted and at worst has the wrong emphasis: either we can give an account of the basing relation, in which case the well-founded solution to the generality problem will work, or we cannot – in which case externalists and internalists alike are embarked in a hopeless task.

ACKNOWLEDGEMENTS

A previous version of this paper was presented at the Fifth Annual Bellingham Summer Philosophy Conference. Many thanks to Michael Bergmann and Andrew Cortens, which were my commentators, as well as the audience, especially John Hawthorne, Baron Reed, Ted Sider and Brian Weatherson for their comments and criticisms. Many thanks also to Earl Conee, Carolina Sartorio, Alan Sidelle and Ernest Sosa.

NOTES

[1] The counterfactual condition is needed in order to account for the possibility of highly reliable but never used processes, on the one hand, and of processes that actually yield lots of true beliefs but are intuitively unreliable because they would very easily yield false beliefs, on the other hand. Much more, of course, is needed in order to fully characterize an epistemically useful notion of reliability. But the problems and issues that will be examined here can be dealt with at this rough level.

[2] Cf. Conee and Feldman (1998). Feldman had presented the problem in Feldman (1995), and Goldman himself (a leading proponent of reliabilism) had first called attention to it in Goldman (1979).

[3] I am not aware of any published defense of this solution to the generality problem, but it is the kind of thing that one hears every once in a while in the hallways.

[4] Alston (1995), p. 2.

[5] Carolina Sartorio and Ted Sider convinced me (in private conversation) that, usually, only a subset of the worlds containing a given token process will be relevant to determining the reliability of that token. To use an example by Sider, the same token process that, on a sunny day, produces in me the belief that there is a daisy in the field could have operated on a rainy day. Intuitively, the truth value of the belief produced by the process on a rainy day is irrelevant to its reliability on a sunny day.

[6] Or, more cautiously: if we have a solution to the problem of selecting the worlds inhabited by the token process that are relevant to determining the reliability of that process, then we have a solution to the generality problem for reliabilism. Whether the problems are the same depends on the (relatively uninteresting) issue of how to count problems.

[7] How does a set of possible worlds containing a process determine a unique type to which the process belongs? As follows: the relevant type is the type having all and only those properties had by the process in all of the possible worlds that belong to the set (if the process has the properties in the actual world, then they are regular properties of the process, otherwise they are possible properties).

[8] For a similar argument that focusing on token processes will not help the reliabilist, see Mark Heller (1995), p. 514, n. 7.

[9] Although the fact that an epistemic theory is externalist clearly doesn't entail that it is a theory of doxastic justification. For instance, an externalist theory could say that a proposition is justified for any subject if and only if the objective chance of its being true is high.

[10] Although, again, the fact that an epistemic theory is internalist clearly doesn't entail that it is a theory of propositional justification. For instance, an internalist theory could say that a subject is justified in actually believing a proposition if and only if that subject is completely convinced that the proposition is true.

[11] Or, at least, adequate theories of doxastic justification do. One could, of course, cook up a theory of doxastic justification doesn't incorporate the idea of well-foundedness – for an example, see the internalist theory of doxastic justification that I construe in note 10.

[12] Feldman and Conee (1985), p.15.

[13] Feldman and Conee recognize this in note 2 of Feldman and Conee (1985). It is worth noting that the proper clarification of either the notion of "fit" or the idea of "having" some evidence might appeal to external factors, which would obviously defeat the intended internalist character of the theory. Indeed, my proposal about how to deal with the generality problem can be seen as an evidentialist characterization of well-foundedness which explains the notion of fit by appealing to reliability.

[14] It is important to note that I have argued that every adequate epistemological theory will have to appeal to the basing relation, but I haven't argued that for any specific analysis of that relation. In particular,

although I think that the basing relation should be causal, the points that I make in this paper are independent of this claim.

[15] Alston (1995), p. 17.

[16] See Conee and Feldman (1998), Adler and Levin (2002) and Feldman and Conee (2002).

[17] Conee and Feldman (1998), p. 12.

[18] Adler and Levine (2002), p. 92.

[19] Feldman and Conee (2002), p. 100.

[20] When the process in question takes beliefs as inputs, we have to take into account the distinction between unconditional and conditional reliability – cf. Alvin Goldman (1979).

[21] See Stewart Cohen (1984).

[22] For details on this solution, see Comesaña (2002). I add two brief notes to the discussion of condition (ii). First, Earl Conee commented (in personal communication) that there might be other cases that can be constructed where the type identified by well-founded reliabilism is reliable because there is a necessary connection between the evidence and the truth of the belief formed on that basis, and it might be the case that, on some of those cases, the belief is not intuitively justified. I agree that necessary connections between evidence and propositions can cause trouble for reliabilism. Ultimately, I think that the way out of this problem is to require reliability only as a necessary condition for justification, and not a sufficient one – to replace the "if and only if" in the definition of well-founded reliabilism with an "only if." This will give us something less than a full theory of epistemic justification, but the resulting theory will (a) still be a recognizable form of externalism, and (b) will still provide a solution to the generality problem.

Second, as Michael Bergmann suggested in his comments on this paper at Bellingham, there is something strange in claiming that a reliabilist theory would not have solved the generality problem unless it is materially adequate – which is the effect of condition (ii). The new evil demon problem and the problem of propositions whose truth is necessarily connected to some piece of evidence seem *prima facie* independent of the generality problem, and it seems therefore unfair to criticize a proposal about how to solve one of them by noticing that it doesn't solve the others. (Of course, a proposal to solve one of the problems shouldn't make it impossible to solve the others, but that is not in question here.)

[23] Conee and Feldman (1998), p. 5.

[24] See note 13.

[25] Thanks to Alan Sidelle and Ernest Sosa for bringing this worry to my attention.

[26] And appealing to clause (ii) (c) is not going to help, because the case can be told so that it *is* satisfied – for instance, by making sure that the

evidence that favors the proposition that Mr. Sloppy is unreliable in the current circumstances is not available to Mr. Sloppy.

[27] Although the following idea is not *completely* awful: that, no matter how the facts turn out, it will always be possible to correctly describe them by saying that someone with superior discriminatory capacities has the ability to collect better evidence.

[28] Conee and Feldman, (1998), p. 24.

REFERENCES

Adler, J. and Levin, M. (2002): 'Is the Generality Problem too General?', *Philosophy and phenomenological Research* LXV(1), 87–97.

Alston, W. (1995): 'How to Think About Reliability', *Philosophical Topics* 23, 1–29.

Cohen, S. (1984): 'Justification and Truth', *Philosophical Studies* 46, 279–295.

Comesaña, J. (2002): 'The Diagonal and the Demon', *Philosophical Studies* 110, 249–266.

Conee, E. and Feldman, R. (1998): 'The Generality Problem for Reliabilism', *Philosophical Studies* 89, 1–29.

Feldman, R. (1995): 'Reliability and Justification', *The Monist* 68, 159–174.

Feldman, R. and Conee, E. (1985): 'Evidentialism', *Philosophical Studies* 48, 15–34.

Goldman, A. (1979). 'What is Jusitified Belief?', in G. Pappas (ed.), *Justification and Knowledge* (pp. 1–23), Dordrecht: D. Reidel.

Heller, M. (1995): 'The Simple Solution to the Problem of Generality', *Nôus* 29, 501–515.

Feldman, R. and Conee, E. (2002): 'Typing Problems', *Philosophy and Phenomenological Research* LXV(1), 98–105.

Department of Philosophy
University of Wisconsin - Madison
5163 Helen C. White Hall
600 North Park Street
Madison, WI 53706
USA
E-mail: jmcomesana@wisc.edu