E. CONEE and R. FELDMAN

THE GENERALITY PROBLEM FOR RELIABILISM

(Received 8 October 1997)

I. INTRODUCTION

Reliabilism and the Generality Problem

Reliabilism is the most widely discussed contemporary epistemological theory. The most widely discussed version of reliabilism is process reliabilism, which makes the processes that cause and sustain beliefs epistemically crucial. The central idea of process reliability theories of epistemic justification is this:

RJ. A belief is justified if and only if it is produced by a process that reliably leads to true beliefs.[1]

A fully articulated reliabilist theory must identify with sufficient clarity the nature of the processes it invokes. In doing so, the theory confronts what has come to be known as 'the generality problem'.[2]

A simple example will show the nature of the problem. Suppose that Smith has good vision and is familiar with the visible differences among common species of trees. Smith looks out a house window one sunny afternoon and sees a plainly visible nearby maple tree. She forms the belief that there is a maple tree near the house. Assuming everything else in the example is normal, this belief is justified and Smith knows that there is a maple tree near the house. Process reliabilist theories reach the right verdict about this case only if it is true that the process that caused Smith's belief is reliable. And one might think that the process is obviously reliable. However, before accepting this conclusion, we should think carefully about exactly what that process is and what its reliability consists in.

Light reflects from the tree and its surroundings into Smith's eyes. Optic neural events result, and these produce further neural events within Smith's brain. Particular concrete occurrences, involving sensory neural simulation in combination with complex standing con-

Philosophical Studies **89:** 1–29, 1998.

ditions in Smith's brain, result in Smith forming the belief. This sequence of concrete events is the process that caused the belief. So, if we take the process that must be reliable to be composed of causally active events that bring about the belief, then reliabilism requires for justification that a sequence of concrete events is reliable.

However, reliability is a kind of tendency. The notion of reliability applies straightforwardly only to enduring mechanisms, such as an eye or a whole visual system, and to repeatable types of processes, such as the type: visually initiated belief formation. Reliability does not apply in any obvious way to the particular sequence of concrete events that caused Smith's belief on this occasion. Each event in the sequence happens only once and the sequence causes whatever beliefs result just on that occasion. Process reliabilists who realize this have sought the requisite reliability in the *types* of process of which particular causal sequences are tokens.[3]

As many reliabilists have recognized, each token process that causes a particular belief is of numerous different types of widely varying reliability. The token event sequence in our example of seeing the maple tree is an instance of the following types, among others: visually initiated belief-forming process, process of a retinal image of such-and-such specific characteristics leading to a belief that there is a maple tree nearby, process of relying on a leaf shape to form a tree-classifying judgment, perceptual process of classifying by species a tree located behind a solid obstruction, etc. The number of types is unlimited. They are as numerous as the properties had by the belief-forming process. Thus, process reliability theories confront the question of which type must be reliable for the resulting belief to be justified. It is clear that the answer to this question will significantly affect the implications of the theory. For instance, while visually formed beliefs in general seem to be fairly reliable, processes that use a characteristically maple-leafish visual experience to judge that a maple tree is near seem much more highly reliable, and perceptual processes leading to a belief that a tree, which is behind a solid obstruction, is of a particular species seem generally unreliable, in spite of the fact that in some of their instances, such as the present case, the obstruction is transparent. The process token

is of endlessly many other types as well, types of extremely varied reliability. So, which type has to be sufficiently reliable?

Process reliabilists must solve this generality problem. A solution identifies the type whose reliability determines whether a process token yields justification.[4] This type is "the relevant type" for that token. Thus, it is not the causally active process token that has to be sufficiently reliable, according to reliabilists. It is the relevant type of the process. We need to know what determines this sort of relevance.

Without a specification of the relevant type, process reliabilism is radically incomplete. Only when a bearer of reliability has been identified does the theory have any implications about the justification of beliefs in particular cases. Philosophers often overlook this. They purport to determine whether or not a given belief is justified according to reliabilism using nothing more than one description of the process causing the belief. No such inference is acceptable. The theory must first be elaborated at least enough to imply exactly what process type has to be reliable in the case in question. A fully general reliabilist theory of justification has to do this for all cases in which there is a fact of the matter.

A second necessary task for process reliabilists is to specify which situations of a process type's operation determine whether or not the type is reliable. Strength of reliability might be settled by the frequency with which the process actually produces true beliefs or rather by its truth-to-falsehood output ratio in certain counterfactual circumstances. The generality problem arises no matter how this question about reliability is answered. William Alston's sensible specification of what determines the reliability of a process type will do for present purposes:

> R. A process type is reliable if and only if it would yield a high proportion of truths over a wide range of situations of the sort we typically encounter.[5]

B. Necessary Conditions for a Solution to the Generality Problem

A solution to the generality problem must meet the following three conditions.

First, it must be principled. Given the multiplicity of belief-forming process types and their variations in reliability, it is easy

to make *ad hoc* case-by-case selections of types that match our intuitions. But case-by-case selections of relevant types does not constitute working out a reliabilist theory of justification.

The claim that the reliability of "the relevant type" of the belief-forming process is what determines a belief's justification is analogous to the claim that "the suitable type" of a horse is what produces victory in a horse race. In the absence of further explanation, this use of 'suitable' has no definite content. On its own, the phrase 'the suitable type of horse' tells us nothing about what makes horses win races. If there is no further explanation but rather we are offered case-by-case choices of "suitability-making properties", choices made on the basis of knowing which horses are the winners, then the claim is no closer to having any definite content. Clearly, a general basis for identifying suitability is required for the claim to say more than just that something or other makes each winning horse win its race. Analogously, we have an informative reliabilist theory of knowledge or justification only after we are told what determines "the relevant type" in general.

Although a solution must be principled, it need not state necessary and sufficient conditions for relevance that are either precise or always determinate. Claims to the effect that a belief is "epistemically justified" might be vague and they might be context-sensitive in various ways. A solution must be universal only in that it must specify the relevant type whenever there are definite facts about justification.

The second requirement for solving the generality problem is that the rule must make defensible epistemic classifications. Stating a general rule of relevance that merely assigns some type or other to each process token does not constitute an adequate solution to the generality problem.[6] The types identified must have a reliability that is plausibly correlated with the justification of the resulting beliefs.

Finally, a solution must remain true to the spirit of the reliabilist approach. We are addressing process reliability theories.[7] So, the rule of relevance must somehow implement the basic idea that it is the reliability of a process of belief formation, specified in non-epistemic terms, that settles the epistemic status of the belief. Process reliabilists characteristically think that a belief is justified because the workings of the process that produced it (or sustained it) are suf-

ficiently conducive to generating true beliefs. A solution to the generality problem would specify those workings so as to bear out this idea. A solution thus cannot identify the relevant type for a process in a way that merely smuggles a non-reliabilist epistemic evaluation into the characterization of relevant types. For instance, 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).[8] But this theory is only verbally a version of reliabilism. It mentions the processes 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.

C. Our Thesis

Our thesis is that the prospects for a solution to the generality problem for process reliabilism are worse than bleak. We will investigate the merits of approaches exemplified by several recent proposals. There is no significant progress in any of these approaches, singly or in combination. The basic process reliabilist idea just does not pan out.

It is reasonable to look for a solution to the generality problem in three places: common sense, science, and context. Common sense is the likeliest source. As we shall soon see, Alvin Goldman's early account of reliabilism draws much of its initial attraction from the *prima facie* correlation between justified beliefs and beliefs produced by common sense types of processes that are probably reliable. Goldman immediately realized that some refinement of these common sense types is needed, for reasons that we shall illustrate below. But at first glance the thought is appealing that common sense process types like "careful perception", "vivid memory" and the like are reliable. So, it makes sense to pursue the reliabilist idea that these types of process produce justified beliefs because of their reliability. In contrast, common sense belief-forming process types like "guessing" seem to be unreliable and seem to yield unjustified beliefs.

If, as we shall argue, common sense types will not do, then the next likeliest source of relevance is scientific classification. Scientific types of belief-forming processes are types that correspond to the predicates that enter into the laws and explanations of science. We shall next investigate the possibility of solving the generality problem by identifying relevant types with these scientific types.

Another reasonable thought is that different types are relevant to justification in different contexts, just as different comparison classes determine the application of terms like 'small' and 'far' in different contexts and just as different reference classes determine the truth value of probability judgments in different contexts. Thus, we shall consider next the merits of contextualist solutions to the generality problem.

We shall argue that none of these approaches works out. This might raise the concern that our way of posing of the generality problem for reliabilism is somehow ill-conceived. It might be thought that the relevant types are obvious when the question is properly understood, or that no general solution is actually needed. We shall take up this line of thinking as well.

That exhausts the reasonable philosophical approaches to the generality problem. If they all fail, then so does process reliabilism.

II. COMMON SENSE TYPES

In his pioneering defense of process reliabilism, Alvin Goldman appeals to common sense process types in an effort to convey the plausibility of the theory. He writes,

> . . . what kinds of cause confer justifiedness? We can gain some insight into this problem by reviewing some faulty processes of belief-formation, i.e., processes whose belief-outputs would be classed as unjustified. Here are some examples: confused reasoning, wishful thinking, reliance on emotional attachment, mere hunch or guesswork, and hasty generalization. What do these faulty processes have in common? They share the feature of *unreliability*: they tend to produce *error* a large proportion of the time. By contrast, which species of belief-forming (or belief-sustaining) processes are intuitively justification-conferring? They include standard perceptual processes, remembering, good reasoning, and introspection. What these processes seem to have in common is *reliability*.[9]

Thinking of reliabilism in terms of these types gives the theory its initial appeal.

However, common sense types have two liabilities as the basis for a solution to the generality problem. First, there are far too many common sense types to provide a unique identification of the relevant type for each process token. In our initial example, Smith's visually formed maple tree belief results from a process instantiating all the following common sense types: visual process, perceptual process, tree-identifying process, daytime process, indoor process, etc., etc. These types differ widely in their reliability. So, we still need to be told which one determines the justificatory status of the resulting belief.

The other main problem with the types Goldman mentions is that not all beliefs resulting from any one such type are even approximately equally justified. Consider another common sense type that Goldman refers to, brief and hasty scanning. Sometimes, on the basis of a brief and hasty scanning we can get extremely well justified beliefs, as when we see in a glance that there is a tree in the backyard. Other times brief and hasty scanning does not yield a justified belief, as when the belief concerns exactly how many leaves are on the tree. Simple common sense classifications are thus too broad to make the right epistemic distinctions among beliefs.

In a recent discussion of the generality problem, William Alston sometimes calls what he proposes as the relevant types "habits" of belief formation.[10] Likewise, Charles Wallis appeals to "strategies" of belief formation.[11] Habit and strategy are common sense classifications of some of the ways we form beliefs. In classifying trees by species, an expert naturalist has identifying routines that differ considerably from those of novice and ill-informed tree classifiers, even though all of them may judge by experiencing the same views of the trees. The expert is better justified. So there is some initial plausibility in the idea that it is the "routine", the "habit of mind", the "strategy", employed in forming a given belief that determines its level of justification.[12] This suggests:

H. The relevant type for any belief-forming process token is the habit of mind, or belief-forming strategy, that it instantiates.

For a large class of cases, it is doubtful that (H) serves to identify a single relevant type. This is because many process tokens are instances of more than one habit. Smith, our maple tree identifier, may have a habit of concentrating while making careful visual judgments, a habit of calling to mind types of trees known to be in the area when making species classifications, and a habit of counting points on leaves for identifying deciduous trees. Some of her belief-forming process tokens result from the employment of all three habits. So there would be no such thing as "the habit" employed on those occasions, and thus no relevant type by the present proposal.

There are, furthermore, cases in which justified beliefs are formed in a way that is in no intuitive sense "habitual", or "routine", or "strategy-employing". For instance, Smith might happen to notice a cardinal on a branch of the maple tree, and be thereby justified in believing that a cardinal is there. She is not employing any strategy, or habit, or routine, in forming this belief. Thus, a theory that requires a high enough reliability for the relevant type here would conclude that the belief is not justified, since there is no habit or strategy that is either reliable or unreliable. Yet in many such cases the belief clearly is justified.

Also, the same belief-forming habit can produce some justified beliefs and some unjustified ones. Jones might make a habit of judging the theme of a philosophy article by reading only its concluding paragraph. Sometimes the theme is clearly presented there and Jones will be justified. Other times the final paragraph does not make clear the point of the paper and Jones will not gain justification by employing this procedure.

Another approach using common sense classifications would be to hold that the solution to the generality problem is to classify together processes that produce equally general beliefs:

> G. Two process tokens are of the same relevant type if and only if they generate beliefs at the same level of generality.

(G) has no promise as a solution to the generality problem. The problem of finding the relevant type does not reduce to that of finding the right level of generality for the contents of the resulting beliefs. It is often not clear what "level of generality" a belief has. But

if there is any merit in the approach that (G) represents, then two judgments will be at the same level of generality if their contents consist in classifying an individual by species. Thus, the following visually based beliefs are all at the same level of generality: this is a mountain-goat, this is a giraffe, this is a crocodile, this is an alligator. (G) implies that all such classificatory beliefs result from the same relevant type, and hence all are equally justified. But clearly this is not so. For instance, some such beliefs are based on more justifying perceptible features than others. To ordinary observers, nearby giraffes are pretty obvious, while nearby crocodiles are easily mistaken for alligators. Processes generating equally general beliefs are not all equally justifying.

A similar idea would be to distinguish processes in terms of the identity of their particular output beliefs, so that the different beliefs just mentioned would result from different relevant types of processes. This has numerous unacceptable results too, however. Clearly there are both justified and unjustified examples of belief in the same proposition.

There is no reason to think that any appeal to simple common sense types will solve the generality problem. Their main liability is that they are too broad to differentiate properly among the justification levels of our various beliefs. Less simple types can be constructed by conjoining together the broad common sense classifications that we have been discussing. These can be much narrower, for instance: visual process causing a belief that classifies by species a close, unobstructed, opaque object, in bright sunlight. But the members of such types still vary in their degree of justification depending on such things as whether the viewer is familiar with the visual appearance of the species from the viewing angle, has normal vision, is intoxicated, is expecting visual trickery, is emotionally distraught, etc.. There is no good reason to believe that even such narrow kinds will include only equally justified beliefs, however elaborately they are specified, as long as they use only common sense nonepistemic categories.[13]

Common sense types thus do not stand scrutiny as candidates to provide a satisfactory solution to the generality problem.

III. SCIENCE

It is in keeping with the "naturalistic" spirit of reliabilist theories to look for classifying help from natural science. One tempting line of thought is that reliabilists can count on cognitive psychology to identify the types of belief-forming processes that will be useful to their theory. Suggestions of a such a view can be found in writings by Alston, Goldman and Ralph Baergen.[14]

A. Natural Kinds

Alston's mention of habits of mind is not his theoretical proposal for coping with the generality problem. Rather, he suggests that belief-forming process tokens belong to natural kinds and that these kinds are the types to which reliabilists ought appeal. He writes:

> With a process token, as with any other particular, any of its properties can be said to be correlated with a type to which it belongs . . . Even if it is true that you and I belong to indefinitely many classes, such as *objects weighing more than ten pounds, objects that exist in the twentieth century, objects mentioned in this paper*, etc. etc., it is still the case that membership in the class of human beings is fundamental for what we are in a way that those others are not, just because it is the natural kind to which we belong. I shall suggest that something analogous is true of belief-forming processes – that there are fundamental considerations that mark out, for each such process token, a type that is something like its "natural kind".[15]

Although this is not Alston's final account of the matter, it is important to see that more is needed. Merely citing the fact that each belief-forming process falls into a natural kind does not provide an adequate rule of relevance. To see this, note the inadequacy of the following solution to the generality problem.

> NS1. The relevant type for any belief forming process token is the natural kind to which it belongs.

Process tokens may belong to natural kinds. Still, there is no good reason to think that each token belongs to just a single natural kind, and hence no reason to think that (NS1) provides a solution to the generality problem. What the natural kinds of belief-forming processes are is up for grabs, but every belief-forming process token is categorized in multiple ways by laws in each of several sciences. These all seem to be natural kinds of the process, according to current

science. Reasonable candidates for natural kinds of a typical visual belief-forming process include electrochemical process, organic process, perceptual process, visual process, and facial-recognition process. All belief-forming process tokens are thus in a multiplicity of natural kinds. So (NS1) does not single out a relevant type for any such process. These natural kinds differ widely in their reliability. So, (NS1) does not solve the generality problem.

B. Psychological Realism

Process tokens thus belong to numerous natural kinds. Alston contends, however, that for each belief-forming process token there is only one type that is "psychologically real". His suggestion is that this type is the relevant type.

According to Alston, every process token instantiates what he calls a "function". He stipulates that this term is to have its mathematical sense. In the case of beliefs formed on the basis of perceptual experience, these functions take as inputs features of experience to which we are responsive and yield beliefs as outputs. Alston is aware that each particular input/output pair is in the extension of many mathematical functions, but he claims that there is only one such function that any belief-forming process actually is an "activation" of. Only this one is psychologically real.[16]

The intended solution to the generality problem seems to be:

> NS2. The relevant type for any process token is the natural psychological kind corresponding to the function that is actually operative in the formation of the belief.

(NS2) does narrow the set of candidates for relevant types. Furthermore, psychology does aspire to provide psychological explanations of at least all normally acquired beliefs.[17] If this aspiration is met, there will be psychological types of belief-forming process for all such beliefs.

If (NS2) provides a solution to the generality problem, it must be that there is only one actually operative "psychologically real" type for each belief-forming process. In apparent support of this, while discussing the application of (N2) to beliefs resulting from vision, Alston emphasizes that there is a fact about which elements of a

visual scene a person responds to in forming a belief about what is present. Thus, in our example about Smith and the maple tree, Smith might form her belief on the basis of noticing certain features of leaf shape. The token process therefore goes from these input features to that belief. In other examples, when presented with the same scene Smith might pick up on features such as the tree's overall shape or bark texture, rather than leaf shape. These considerations show that the relevant type in the original case must be one that corresponds to a function having as an input/output pair the leaf-shape features to which Smith responds and the belief that she forms.

This may limit somewhat the candidates for relevant types, but in Smith's case there still are numerous overlapping functional relations, and corresponding psychological process types, that include the input/output pair we've identified. 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 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.

An example from another domain may help to make this point clearer. Suppose that a certain pot of water at sea level is brought to a boil. There occurred a certain sequence of concrete events leading to the boiling of the water. This sequence instantiates any number of types, all "physically real". We can identify these types in terms of the functions that describe their final stage. At any given pressure, there

is a function that maps water onto a certain temperature – its boiling point. This corresponds to the process "bringing water to a boil at sea-level atmospheric pressure". There is a broader type, "bringing water to a boil". The function corresponding to this second type takes water and varying pressures as inputs, and yields a boiling point for water at each temperature. A still broader function takes as inputs triples of temperatures, pressures, and types of liquid and yields the boiling point for each. This corresponds to the type "bringing liquids to a boil". The token process in our example is an instance of all these types. It is not the case that only one is "physically real". All of them accurately characterize what occurred in the pot. Similarly, far too many functions are "psychologically real". They all correspond to natural psychological kinds. So, (NS2) fails to identify the relevant type.

C. Maximum Specificity and Narrow Causal Types

Alston also suggests that his psychological realism implies, or at least is compatible with, a different specification of relevant types, one that relies on completely causally specific functions. He assumes that "the functions in question are maximally specific, in that any difference in input that is registered by the function indicates a different function".[18] Making use of this idea of maximal specificity is one way of trying to make good on the idea that only one function is "operative" in the formation of any belief.[19]

In any case where a person forms a belief on the basis of a perceptual experience, some features of the experience contribute to a belief-forming causal sequence that starts with the experience. Other features of the experience play no causal role. The same goes for subsequent events in the sequence leading to the belief. Some features of these events help to cause the belief, others do not. The maximum specificity proposal is the idea that the relevant type includes all and only process tokens with the same causal features: they all begin with experiences with the same causally active features, are followed by subsequent events with the same causal features, and have the same belief as output. At one time, Alvin Goldman suggested a very similar solution to the generality problem.[20] We can formulate this proposal as follows:

NS3. The relevant type for any belief forming process token *t* is the natural kind that includes all and only those tokens sharing with *t* all the same causally contributory features from the input experience to the resulting belief.[21]

(NS3) does yield a unique type for each process token. But the reliabilist theory of justification that employs (NS3) is seriously defective. (NS3) classifies into the same relevant type only beliefs that share *all* internal causal predecessors. Thus, on the reasonable assumption that the content of any normally formed belief is causally determined by its antecedent psychological causes, according to (NS3) each relevant type can have only one content for its output belief.[22] This makes trouble in cases in which the proposition believed dictates the truth-ratio of all process types leading only to it. In such cases the reliability of the relevant type is settled by the mere identity of the belief. Thus, the relevant type of a process leading to any necessary truth must be perfectly reliable. The relevant type of any process leading to any necessary falsehood must be perfectly unreliable. Also perfectly reliable would be the relevant types of all processes leading to any self-confirming belief, such as the belief that someone believes something. The relevant type of the following beliefs would be perfectly unreliable: the belief that there are no beliefs, and the belief that nothing is caused. Since it seems clear that in all of these cases the beliefs can have a level of justification that is other than the implied extreme, these examples run counter to (NS3).

The problems for reliabilist theories built on (NS3) are not confined to beliefs in necessities, impossibilities, or the relatively unusual beliefs just mentioned. Suppose that Jones looks very carefully at a tree and forms the belief that it is a beech on the basis of seeing features which are in fact distinctive to beech trees. As long as experience of such features happens to help to prompt Jones to believe that it is a beech tree, it does not matter to (NS3) why they do so. It can be for good reasons, for bad reasons, or for no reason at all. Recall (R), which tells us that the reliability of a type is determined by the long run truth ratio of its output when it functions under typical conditions. In the normal worlds used to evaluate the reliability of Jones's tree-identifying process, nothing other than

a beech tree presents Jones with exactly the features that initiate the causal process leading to his belief.[23] This by itself is enough for the theory to imply that Jones's belief is justified, regardless of how much information he happens to have about the look of beech trees. Since the -highly specific causal factors that led to his belief in fact are indicative of only beech trees, his belief must be justified, according to this theory. In the worlds that determine the reliability of the relevant type, only beeches cause the sort of experience that led to his belief that a beech tree is nearby. So this maximally specific type is maximally reliable. Reliabilist theories based on (NS3) thus are unable to distinguish the epistemic status of lucky guesses that happen to be based on distinctive features from expert judgments based on well-understood classifications.

An additional problem is that (NS3) yields a version of reliabilism that is not in keeping with the spirit of process reliabilism. As we have just seen, (NS3) often renders irrelevant the details of the process intervening between an input and a resulting belief. In particular, suppose that Jones and Smith both respond to the same features of a visual input with the belief that there is an elm tree present. Suppose that this input will occur only when there is an elm tree present – it is a distinctive look of an elm leaf, say, the visual appearance of a particular quantity of tiny notches around its edge. Finally, suppose that Smith knows what she is seeing, while Jones is applying some ridiculous and unjustified sort of numerology to the topic. Jones plucks from thin air the idea that the magic number for elms is nine. Jones gets a nine for the tree whose leaf he beholds by counting the number of those distinctive elm notches along the edge of a leaf, and dividing by six, his "tree number". Given (NS3), the relevant types for their processes are maximally specific. These types are thoroughly reliable since nothing other than an elm would cause just that input in any significant fraction of nearby worlds. The fact that one of the two knows what elms look like and the other does not and the fact that one process goes through a silly application of superstitious nonsense do not affect the reliability of the maximally specific types (NS3) specifies.[24] It is just this sort of difference that process reliabilism is supposed to make matter. It is supposed to be sensitive to the possibility that the process one person uses is not generally reliable while the one the other uses is generally reliable,

even if in the case at hand both people happen to begin their processes by noticing what is in fact an extremely reliable indicator of the right answer. In other words, process reliability theories are supposed to appeal to much broader relevant types.

D. Categories from Science

Ralph Baergen discusses several examples, explaining what reliabilists might say about them. By generalizing from his remarks it is possible to devise another way reliabilists might appeal to science to solve the generality problem. It is also a second way of attempting to cash out Alston's remark that only one process type is "actually operative" in belief formation.

One example, discussed in the literature by Richard Feldman, concerns a person who sees something on a distant hill.[25] She forms the belief that what she sees is an animal and the belief that it is a mountain goat. Feldman points out that the more general belief may well be better justified than the more specific one. So, he concludes that reliabilists must find a way to distinguish between the types of processes that cause the beliefs.

Baergen proposes a way to do this.[26] He appeals to David Marr's theory of vision, which holds that in classifying objects on the basis of visual perception, we generate a model of the object which "is compared to descriptions in a sort of catalogue. This catalogue is arranged in levels, so that rough categorizations take place at the lower levels, followed by more fine-grained discriminations at higher levels".[27] Baergen suggests that we make use of this idea in identifying relevant types:

> Our account of processes might well reflect this by saying that rough categorizations are generated by different process[es] than those yielding fine-grained categorizations. Applied to Feldman's case, the mountain-goat belief is generated by a different process than that which generated the animal-belief, for they involve different levels of categorization. Also, the process that generated the animal-belief is likely to be more reliable, for there are likely to be fewer nearby situations in which this generates a false belief than there are for the mountain-goat process. So, Reliabilism *can* provide intuitively correct results here.[28]

No doubt reliabilists can state a rule of relevance that produces the intuitively correct results "here". But reliabilism needs a fully general rule. Baergen reports part of a theory of vision that implies that perceptual classifications result from processes that are organized

by levels of generality of the resulting beliefs. He suggests that reliabilists can identify relevant types in some way that plays on this fact. However, Baergen does not make clear how to build upon this example to develop a general account of relevant types.

One possibility, suggested by Baergen's use of psychology, is that the relevant types are the types that are invoked by the best psychological theories of belief-formation. The idea here is that while any token belongs to numerous types that are psychologically real, only one of those types will enter in the best psychological theory that explains the resulting belief. That type is the relevant type. Marr's theory may have been used to illustrate how this might apply in the case of visual belief-formation.

We can formulate this idea as follows:

> NS4. The relevant type for any belief-forming process token *t* is the psychological kind that is part of the best psychological explanation of the belief that results from *t*.

It may be that Alston had something like (NS4) in mind when he said that only one type was "actually operative".

(NS4) rests on the dubious assumption that there is a unique "best" psychological explanation for each belief. The value of an explanation depends upon the use to which it is put. A very specific and narrow explanation might have greater value for some purposes, while a broader explanation might have greater value for other purposes.[29]

Even if (NS4) did identify unique types, it would not possible to evaluate its implications for process reliabilism without knowing what those types are. There is no good reason to think that the types that are of greatest value for psychological explanation are uniformly helpful to reliabilist theories of justification.

To see why types that are particularly useful for psychological explanation might not be of much help to reliabilists, consider the types Baergen mentions. His proposal ties the relevant types for classificatory beliefs based on visual perception to the level of generality of the resulting belief, and he suggests, plausibly, that a type that produces relatively general beliefs is more reliable than types that produce more specific beliefs. A version of process reliabilism

making use of this idea would thus make more general classificatory beliefs better justified than more specific classifications. That is an unacceptable result. Sometimes, a belief applying a broader classification is less well justified than is a belief applying a narrower one. For instance, Jones might use a visual basis for both his belief that the tree he is near is an elm tree and his belief that the tree he is near is a deciduous tree. He can be less well justified in believing the latter, despite its applying a broader classification. This might be true because Jones does not realize that all elms are deciduous and has just a shaky grip on visual cues to deciduous trees, but he has good training in recognizing elms. Similarly, a person may know at a glance that a thing she sees is a whale, but be less well justified in her belief that it is a mammal. Thus, sometimes the more general belief is the more justified, and sometimes not. So the generality of a visually based classificatory belief does not determine a relevant type that yields a satisfactory version of reliabilism. There is, then, no reason to think that the particular scientific classifications Baergen mentions yield types that are entirely helpful to reliabilism.

Although science does provide the tools to narrow the candidates for relevant types, there is no good reason to think that scientific classifications provide the tools for solving the generality problem.

IV. SOLUTIONS WITHOUT A NECESSARY AND SUFFICIENT CONDITION

Some philosophers have responded to the generality problem by explicitly denying that the problem requires a general resolution. We will examine two such responses in this section.

A. Constraints

Frederick Schmitt proposes five constraints on which process types are relevant, and then appeals to the constraints in describing problem cases.[30] According to Schmitt, "relevant processes are cognitive processes".[31] His constraints require, among other things, that relevant types are salient, that they are folk psychological process types, and that tokens of the same type are intrinsically similar.

These constraints are not meant to compose what Schmitt calls a "criterion of relevance": a necessary and sufficient condition for

relevant types. Schmitt believes that no such criterion is needed. Instead, the constraints are supposed to identify the sorts of factors that we take to matter when we make judgments about justification.

To explain why no criterion of relevance is needed, Schmitt writes:

> [W]e have intuitions about which processes are relevant. In judging whether a subject is justified in an inferential belief, we check to see which inferential process the subject exercises – e.g., whether it is induction from sufficiently many instances or affirming the consequent. We have the intuition that these are the relevant processes to consider. In the case of perceptual belief, we check which environmental conditions obtain – whether it is sunny or foggy – and whether the subject is careful and attentive in perception or quick and distracted. Here again we have intuitions about which processes are relevant. Reliabilism may explain why perceptual or inferential beliefs are justified or unjustified by relying on these intuitions.[32]

The existence of these intuitions does not relieve process reliabilists of the responsibility to provide an explanation of their invocation of relevance. Granting that the intuitions exist, the question that we have been asking remains to be answered: According to reliabilism, which type must be reliable for a particular belief to be justified?

Furthermore, Schmitt is mistaken about exactly what intuitions we do have. Schmitt says that "we have intuitions about which processes are relevant". Since Schmitt is addressing the generality problem, this claim seems intended to imply that 'relevant' in the reliabilist use of the 'the relevant type of the process' has some intuitive application to examples. But that is not so. The reliabilist use of 'the relevant type' is entirely technical. The expression might as well have been 'the type that determines justification according to the philosophical theory known as 'reliabilism'. No one has pre-analytic intuitions about this topic. It is up to reliability theorists to assign reference to the term from scratch.

Philosophers and others do make intuitive judgments about which features in examples are "relevant" to the justificatory status of beliefs. Schmitt is entirely right to say that in evaluating inferential beliefs we are inclined to judge relevant the pattern of inference followed, and in evaluating perceptual beliefs we judge the environmental conditions and attentiveness of the perceiver to be relevant. We also judge to be relevant the quality and quantity of evidence the believer has. We typically judge to be irrelevant the day of the

week on which the belief is formed and the color of the believers socks. These are not intuitions about which process types are relevant. They are intuitions directly about what determines a belief's epistemic justification.

The existence of intuitions about which factors are relevant to justification does not eliminate reliabilism's need for a theory of relevant types. The constraints Schmitt describes do not do this on their own. They provide a variety of conflicting criteria. In his discussion of cases, Schmitt gives the constraints differing weights so as to achieve the desired result.[33] Perhaps one can, by weighing one factor heavily in one case, a different factor heavily in another, manipulate the constraints in a way that seems to give reliabilism acceptable results. But this is no victory for reliabilism. One could equally well say that the justification of a belief is a function of epistemically irrelevant factors such as the duration of the token of the cognitive process that caused it, the distance of the proximate external cause of the process from the center of the earth, and the amount of energy the process consumed. By *ad hoc* weighting of these factors, one could get acceptable results. The theory, nevertheless, has no merit.

A set of flexible constraints does not solve the generality problem. There are, of course, terms in our language whose application is governed by a set of flexible and varying factors. For example, when we say that someone is a "good athlete", there are a variety of factors that enter into our evaluations. They might include speed, strength, and endurance, among other things. But there is no fixed weight uniformly given to these factors. In different contexts these different factors may be weighed differently and it would be a mistake to ask for some fixed ranking of the importance of these various factors in evaluations of athletic ability. Although Schmitt does not say this, it is possible that he intends to propose that evaluations of processes as reliable work in somewhat the same way.[34] We turn in the next section to a proposal along these lines.

B. Context

Mark Heller contends that the demand for "a general principle for selecting the correct level of generality [for relevant types] . . . is unreasonable".[35] He thinks that contextual factors determine relevant

types and thereby solve the generality problem. Heller elaborates his claims about the role of context as follows:

> 'Reliable' is a perfectly ordinary word that in perfectly ordinary situations is applied to tokens which are instances of several types, where those types have different degrees of reliability. Yet we somehow manage to use this word without difficulty in ordinary discourse.[36]

Heller says that the primary task of his paper is to defend the claim that 'reliable' is richly sensitive to the evaluator's context. This much is unobjectionable. The word 'reliable' surely is context sensitive. That is, whether or not a thing is accurately called 'reliable' depends in part upon the standards set by the context of the ascription. These standards vary, depending for instance on how important it is to rely on the thing that is said to be reliable. This is at most a first step toward solving the generality problem. We need to see how context sensitivity helps with the identification of the relevant type.[37]

Heller does not claim just that the standards for the application of 'reliable' are context dependent. He makes the further claim is that we readily understand applications of 'reliable' to process tokens that are instances of many types. Thus, when a person says "that process is reliable", the person can refer to a process token and say something true. The person's statement is true provided the contextually determined type for that token is truly said to be 'reliable' in the context of attribution. If Heller is right, then context determines two features of our predications of 'reliable' to tokens. One has to do with the standard for the strength of reliability required for the term to apply in the context. That feature is of no help in determining the relevant type. The other feature has to do with the identification of the type that must meet those standards. We will refer to these latter types as 'contextually determined types'. Thus, a phrase of the form 'the process leading to S's belief that p' is supposed to have, relative to a context, a contextually determined type.

A solution to the generality problem can be constructed from these thoughts. The proposal that we shall formulate combines Heller's contentions about the context dependence of the word 'reliable' with the epistemic contextualist view that the standards for assessing the truth value of knowledge and justification attributions is dependent on the attributor's context.[38]

C. In any context, C, if a person says something of the form 'S knows p' or 'S is justified in believing p', the relevant type of the belief-forming process is the contextually determined type for the phrase 'the process leading to S's belief that p' relative to context C.

(C) embodies the idea that the description 'the process leading to S's belief that p' has a contextually determined process type. (C) puts that idea to the service of reliabilism by identifying contextually determined types with the relevant types needed to fill out reliabilist theories of knowledge and justification.

A fundamental objection to (C) is that contextual factors do not typically yield one determinate process type for the phrase 'the process leading to S's belief that p'. As a result, reliabilist theories built upon principle (C) will not yield the correct truth value for many clearly determinate attributions of knowledge or justification.

There are some situations in which phrases referring to process tokens apparently work in the way Heller describes. For example, suppose Jones says, "I have three ways to start my old jalopy: first, shifting into gear while rolling it down a hill; second, jump-starting it; and third, praying and then turning the key. Only the first two usually work." Suppose that Jones then starts his car by jump-starting it. He remarks:

P. "The process by which I just started my car is reliable".

Here, Jones's explicit mention of the three types serves to limit drastically the types under consideration. The token mentioned in (P) is of one of those types only. So, this is a case in which 'reliable' is explicitly predicated of a process token and we have no problem in understanding what type must be reliable for the predication to be true.

In typical knowledge attributions, however, no contextual narrowing of candidate process types occurs. If it did, then when a person said that someone knows something, there would typically be a range of contextually salient process types such that the token process leading to the person's belief instantiated only one. But this is plainly not the case for most knowledge attributions. Ordinarily,

no class of types of belief-forming processes will have been made contextually salient. And nothing else about typical contexts isolates any one type. So, it is just not true that in the context of knowledge attributions there are contextually determined types for the phrase 'the process that caused this belief'.

To see that this is so, consider our initial example in which Smith comes to know that there is a maple tree nearby by seeing it there. Suppose that Jones, who is sitting in the room with Smith, says:

K. "Smith knows that there is a maple tree nearby".

If Heller's version of reliabilism is to work, there must be, relative to the context of Jones' remark, some contextually determined type for the phrase 'the process that caused Smith's belief'. What would that type be? Nothing beyond the speaker's intentions seems to narrow the candidate pool in this sort of example. Perhaps Jones would be thinking of something like perception of familiar objects at a reasonable distance, or perhaps to something narrower, such as visual perception of familiar well-lit trees from a reasonable distance. Perhaps Jones would not have any type of belief-forming process in mind. After all, he did not say anything about belief-forming processes and there is no reason to think that he was having any thoughts about them. So, there is no reason to think that in this sort of mundane example, there is such a thing as the contextually determined type for the phrase 'the process that caused Smith's belief'. Moreover, there is no reason to think that the truth value of Jones's attribution of knowledge to Smith depends in any way on which, if any, of these types Jones has in mind.

Furthermore, even if an attributor of knowledge does have some belief-forming process types in mind, the attributor's thoughts do not identify relevant types in a way that is uniformly helpful to reliabilists. An attributor of knowledge may be mistaken about the reasons for a person's belief, and thus may be thinking about process types that the subject's token process doesn't even exemplify. For example, suppose that Jones witnesses Smith identify a bird as being of a certain species after Smith has had only the briefest glimpse of it under poor lighting conditions. Jones says that Smith's belief is unjustified and so Smith lacks knowledge. Jones does have in

mind some process type for Smith's belief, something like forming a bird classifying belief on the basis of a brief glimpse in poor lighting conditions. Suppose, however, that Smith has formed her belief on the basis of hearing the bird's song, an identification method that Jones has not even thought of. Moreover, Smith does have knowledge as a result. If process reliabilism is anywhere close to the truth about knowledge and justification, it is the reliability of some process type that Smith actually underwent that matters here. So, the generality problem must be solved by appeal to facts about the processes actually involved in the formation of the belief, not by appeal to the possibly mistaken thoughts about those processes in the minds of knowledge attributors.

(C) is incorrect. There simply are no contextually determined types in many, perhaps most, typical contexts in which knowledge and justification claims have a clear truth value. It is true that context helps to determine the standards a process type must meet to be correctly described as "reliable". But the attributor's context comes nowhere near to picking out a relevant type of each belief-forming process, and the process types that are salient to the attributor can be entirely irrelevant to the truth of knowledge claims.

This section has focused on common sense types of belief-forming processes. There are also the many scientific types that classify each belief-forming process. It is clear that nothing about typical contexts of belief, or typical contexts of attribution of knowledge or justification, uniformly singles out one of them. Since our minds are rarely scientifically orientated, speakers' intentions are even less likely to narrow down the scientific types. Nothing else about a context of utterance does so either. Thus, context does not solve the generality problem.

VI. CONCLUSION

That is the full variety of existing approaches to disposing of 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.[39]

NOTES

[1] Some authors discuss process reliability accounts of knowledge rather than accounts of epistemic justification. No point will be made below that turns on the differences between knowledge and justification.

[2] Alvin Goldman in "What is Justified Belief?" in G.S. Pappas (ed.) *Justification and Knowledge* (Dordrecht, Holland, 1979) and *Epistemology and Cognition* (Cambridge, MA, Harvard University Press, 1986) defends process reliabilist accounts of epistemic justification. In those works he recognizes the existence of the generality problem. See especially "What is Justified Belief?", p. 11 and *Epistemology and Cognition*, pp. 49–51. The problem is emphasized in Richard Feldman's "Reliability and Justification", *The Monist* 68 (1985): 159–174. It is also discussed by John Pollock in "Reliability and Justified Belief", *Canadian Journal of Philosophy* 14 (1984): 103–114. For responses to the problem, see the works of William Alston, Ralph Baergen, Mark Heller, Frederick Schmitt, Ernest Sosa, and Charles Wallis cited and discussed below.

[3] It is possible to construct a version of process reliabilism which is only about process tokens and does not confront the generality problem. It faces a considerable problem in making sense of the claim that a token sequence of events has some tendency toward producing beliefs whose truth-ratio would constitute its "reliability". Furthermore, the problems that affect (NS3) below, in virtue of types having just one belief content in their outputs, also affect reliability theories that locate a sort of reliability in process tokens.

[4] There may not always be a fact of the matter. In the examples used here the belief is either definitely justified or definitely unjustified. The reliability of relevant types for process tokens that lead to beliefs whose epistemic status is unclear will be of less value to present concerns, since such cases are less useful in assessing epistemological theories.

[5] "How to Think About Reliability", *Philosophical Topics* (Spring 1995): 1–29. The proposal mentioned here appears on p. 10. If a satisfactory solution to the generality problem existed, it would be worth addressing difficulties with details of this proposal. For one thing, it is not clear who "we" are supposed to be: all of humanity, or all sentient life on earth, or sentient life everywhere in the universe, or etc. And for another thing, it is unclear which belief-forming situations are "typical". Presumably, bizarre psychology lab situations are atypical. But is perception during space travel atypical, no matter how common it becomes? Are situations of fatigue, intoxication, and excitement atypical? Another difficulty is that we may be specially perceptive during rare emergency conditions. If these are atypical situations, then the justified beliefs from these perceptions might turn out not to be of generally reliable types. In any event, if the present work is correct in its main thesis, then these difficulties are not worth pursuing because the generality problem is insoluble.

[6] In some passages in "How to Think about Reliability" Alston seems to construe the generality problem somewhat differently. For his purposes, a solution need only show that there are "objective, psychological facts of the matter that pick out a unique type as the one of which a particular process is a token". (p. 5) Thus, he is content to identify relevant types, leaving as a different matter the question of the acceptability of the resulting reliabilist theory. The problem discussed here is that of getting the theory stated *and getting it right*. Any rule of relevance that

selects one type for each token will generate some reliabilist theory or other, most of them preposterous.

[7] Reliabilist theories that make use of the reliability of indicators or mechanisms of belief-formation are thus not our topic. But the problems for the theory of relevance (NS3) below carry over straightforwardly to many reliable indicator theories. Also, there is a problem similar to the generality problem concerning "the mechanism" that produces a given belief. For instance, when a visual judgment relies on only black-and-white discrimination, is the person's whole visual apparatus the relevant mechanism, or is it the black-and-white sensitive portion of that apparatus, or is it only the active part of that portion? Does "the mechanism" for remembered beliefs include parts of the brain active in forming the belief, or just parts active in storing it and recovering it? These questions may have answers that are attractive to reliabilists, but as with the generality problem, the challenge is to identify a principle that implies all and only the correct answers to such questions.

[8] The results of this theory may be implausible in "demon worlds" in which a demon sees to it that believing in accord with one's evidence does not reliably lead to truths. Whether this is a decisive objection to our evidentialist pseudo-reliabilism depends in part on how reliability is measured. The objection as it is often described makes the challengeable assumption that a process is reliable in a world only if it regularly leads to truths in that world. In contrast, see for instance William Alston's proposal, stated as (R) above. It does not imply that unreliability in a demon world entails a lack of justification. What (R) makes decisive is roughly the truth-ratio of belief-production in more typical situations.

[9] "What is Justified Belief?" pp. 10–11.

[10] "How to Think about Reliability", p. 13f.

[11] Charles Wallis, "Truth-Ratios, Process, Task, and Knowledge", *Synthese* 98 (1994): 243–269. See especially p. 266. Wallis relies on belief-forming strategies as part of his response to problems that he discusses for reliability theories of knowledge. It is not clear that he is attempting to solve the generality problem that is the topic of this essay. One reason for this unclarity is that Wallis is working on a concept of knowledge that is relativized to the specification of a task, unlike the traditional concept which is our topic. In any case, we do not intend to attribute to him a simple reliance on strategies as a full solution.

[12] What follows is a possible solution to the generality problem, suggested by some of Alston's language, that merits a brief look. It is not what Alston proposes. His proposals will be taken up shortly.

[13] In "What is Justified Belief?" Goldman introduced a distinction between belief-dependent belief-forming processes and belief-independent belief-forming processes. The former processes take beliefs, as well as other factors, as inputs and yield new beliefs as outputs. The latter processes do not take prior beliefs as inputs. Belief-dependent processes are reliable when, over a suitable range of cases, they yield true beliefs if their input beliefs are true. Furthermore, a belief resulting from a belief-dependent process is justified only if the input beliefs are themselves justified. One might hope to appeal to this distinction to help deal with some of the examples discussed in this section and elsewhere in this paper. For example, if one's background beliefs are part of the cause of one's animal classifying beliefs, then the differences in the degree of justification for the beliefs mentioned here might be attributable to differences in the degree of justification of

the beliefs upon which they depend. One might therefore be able to maintain the claim that one relevant type is responsible for all the species classifying beliefs.

Defenders of reliabilism have not made significant use of the belief-dependent/belief-independent distinction in their efforts to solve the generality problem. There are good reasons for this. First, it is likely that virtually all beliefs that adult humans form are partially caused by other beliefs. Hence, virtually all our beliefs result from belief-dependent processes. It is therefore doubtful that there is any acceptable way for reliabilists to account for the differing epistemic status of the background beliefs in the examples under discussion. Furthermore, some account of the reliable types for belief-dependent processes is needed. If they are identified in terms of, say, patterns of inference, then process reliabilism turns out to be equivalent to the view that a belief is justified if it results from an inference that is likely to be truth preserving from justified beliefs. This familiar view violates the spirit of process reliabilism since it uses processes only as an indirect way to refer to inferential relations. Finally, it is difficult to see just how to make use of the belief-dependent/belief-independent distinction in conjunction with the specific proposals discussed here. Consider, for example, (G). According to (G), the relevant type is determined by the level of generality of the resulting belief. Thus, according to (G), if two people end up believing that there is a giraffe nearby, they have used processes of the same relevant type. None of the details of the routes by which they got to that belief play any role in determining which type they used. One could be making an invalid inference from justified premises while the other is making an accurate classification based on background knowledge. A theory employing (G) incorrectly evaluates the two beliefs the same way.

As the solutions proposed in the existing literature are discussed below, the reader is invited to note that, like (G), they do not give any role the difference between belief-dependent and belief-independent processes.

[14] Alston's and Baergen's implementations of this idea are discussed below. Goldman mentions this sort of approach in *Epistemology and Cognition*, p. 50.

[15] "How to Think about Reliability", p. 11.

[16] "How to Think about Reliability", section vi.

[17] Philosophers often invoke examples in which beliefs result from blows to the head or tumors. It may be that such beliefs do not result from any *psychological* belief-forming process type. Perhaps the explanations of such beliefs must come from a different science or perhaps psychology must be inclusive enough to account for them too, simply because they are mental effects. If some beliefs lack any psychological cause, that would present a problem for (NS2), since even these beliefs can be assessed for justification, and hence they must have a relevant type.

[18] "How to Think about Reliability", p. 26.

[19] Throughout this section, when we speak of maximally specific functions or types, we mean the maximally specific *psychological* functions or types.

[20] *Epistemology and Cognition*, p. 50.

[21] Theories can differ over exactly what counts as the input. The process type could begin at the surface of the skin, or farther in at some point where conscious experience begins, or farther out in an external cause of the experience. Alston favors perceptual experiences as the initial step (pp. 12f). He does not defend this selection. No point made here depends on any particular beginning for the causal sequence that constitutes the process.

[22] Strictly speaking, the assumption may imply only that the "narrow" content of the beliefs resulting from a given relevant type will be the same. No point made here depends on the difference between narrow and broad content. Also, see note 17 above concerning the completeness of psychological explanation.

[23] It is safe to assume that many of our clear vivid experiences of complex ordinary things like trees are produced only by these same ordinary things in all situations of the sort we typically encounter. Holograms, hallucinations, and perfect pictures are, at most, highly atypical.

[24] One might think that the fact that Jones relies on unjustified background beliefs has some bearing on this example. That thought seems right. But (NS3) ignores this fact and suggests nothing about how to make use of it in defending a process reliabilist theory. See note 13.

[25] "Reliability and Justification", *The Monist* 68 (1985): 159–174. The example is discussed on pp. 164f.

[26] Ralph Baergen, *Contemporary Epistemology*, (Harcourt Brace, Fort Worth, 1995), p. 99. Contrary to what Baergen says, Feldman does not assert that the processes are of the same type. He merely points out the undesirable consequence of the proposition that they are of the same type. It is notable that this sort of example shows that common sense process types, like the visual belief forming process, do not produce beliefs of equal justification even when relativized to a fully detailed specification of the external circumstances.

[27] *Contemporary Epistemology*, p. 100.

[28] *Contemporary Epistemology*, p. 100.

[29] Compare the water boiling example above. There seems to be no reason to think that the explanation at one level of generality is necessarily better than an explanation at any other level.

[30] *Knowledge and Belief*, (Routledge, Chapman, and Hall, New York, 1992), Chapter VI.

[31] *Knowledge and Belief*, p. 169.

[32] *Knowledge and Belief*, pp. 141–2.

[33] For example, Schmitt says about an example that one constraint, which favors a broad relevant type, outweighs two others that favor a narrower type (p. 171). In another case, the existence of two constraints favoring a narrower type is said to outweigh one pointing in a different direction (p. 157).

[34] Schmitt does say that relevance is a "messy, more contextual affair" than some might think (p. 159).

[35] "The Simple Solution to the Problem of Generality",*Nous* 29 (1995): 501–515. The quotation is from p. 502.

[36] "The Simple Solution to the Problem of Generality", p. 502.

[37] Ernest Sosa suggests a contextualist response to the generality problem in *Knowledge in Perspective: Selected Essays in Epistemology*, (Cambridge, Cambridge University Press, 1991). Sosa suggests in a programmatic way that relevant types are ones that can "be usefully generalized upon by us as the epistemic community of the" believer (p. 284). Sosa does not elaborate upon this idea, which is a small part of a complex theory. What he does say does not seem to identify a unique type, since multiple types may be "usefully generalized upon".

[38] Although the following thesis is suggested by much of what Heller writes, it goes beyond the explicit proposals in Heller's paper. Also, it makes no use of passages suggesting that a relevant alternatives approach to a theory of knowledge

solves the generality problem. We see no plausibility in this latter suggestion on its own, and no way incorporate into it the central theme of Heller's paper concerning the importance for solving the generality problem of the context-sensitivity of 'reliable'.

[39] An earlier version of this paper was presented at a symposium at the American Philosophical Association in December, 1996. We are grateful to William Alston for his comments. We are also grateful to Ralph Baergen and John Bennett for comments on earlier drafts.

Department of Philosophy
University of Rochester
NY 14627-0078, Rochester
U.S.A.