NOÛS 38:1 (2004) 177–195

The Generality Problem, Statistical Relevance and the Tri-Level Hypothesis

JAMES R. BEEBE
Louisiana State University

I. Introduction to the Generality Problem

Reliabilism has been the most widely discussed theory of justified belief of the last two decades. It has not, however, been the most widely accepted theory. According to the reliabilist analysis of justified belief,

R1) *S*'s belief in *p* is justified iff it is caused (or causally sustained) by a reliable cognitive process, or a history of reliable processes.[1] (Goldman 1994, p. 309)

A cognitive process will be reliable just when it yields a sufficiently high ratio of true to false beliefs. If a belief is produced by a process with a high degree of reliability, then that belief will have a high degree of justification. If, however, a belief is produced by a cognitive process with a low degree of reliability, then that belief will have a low degree of justification.

After two decades of debate, a few objections have emerged as the standard objections to reliabilism. The Generality Problem is one such objection, the most visible proponent of which has been Richard Feldman (1985; Conee & Feldman 1998). It is now cited as a serious problem for reliabilism in almost every introductory text on epistemology.[2] In this article I offer a solution to the Generality Problem.

The Generality Problem arises because reliabilists claim that it is process *types* rather than process *tokens* that are the bearers of reliability. A process token is an unrepeatable, causal sequence occurring at a particular time and place. Consequently, you cannot ask whether a process token is reliable (i.e., whether it would produce mostly true beliefs over a wide range of cases). Accordingly, reliabilists have claimed that only process types can be reliable or unreliable. We can revise (R1) to take this point into account.

R2) *S*'s belief in *p* is justified iff it is caused (or causally sustained) by a token process that belongs to a reliable cognitive process type (or by a history of such tokens, each of which belongs to a reliable cognitive process type).

It is precisely at the point of distinguishing process types from tokens that the Generality Problem rears its head. As Earl Conee and Richard Feldman (1998, p. 2) put it, "each token process that causes a particular belief is of numerous different types of widely varying reliability." If a belief's justification is a direct function of the reliability of the process type that produced that belief, it is important that there be *only one* of these numerous process types that determines the belief's justification.

Suppose that there could be more than one relevant process type for a given process token. If those process types differed in their degree of reliability, the belief token produced by the target process token might be both justified (in light of the reliability of one process type) and unjustified (in light of the unreliability of some other process type). Since reliabilists claim that a belief is either objectively justified or objectively unjustified, this result would be unacceptable. Reliabilists are committed to there being a single relevant process type for each process token that issues in a belief token. What makes a particular process type relevant? Proponents of the Generality Problem suggest that reliabilists have no principled way to answer this question.

To illustrate the difficulty, Conee and Feldman offer the following story: suppose that Smith looks out his window, sees a maple tree, and forms the belief that there is a maple tree nearby. If perceptual conditions are normal (i.e., Smith has normal eyesight, is not hallucinating; there is adequate sunlight, no occluding objects, etc.), Smith's belief is plausibly taken to be justified. However, the token process responsible for Smith's belief is a member of (at least) the following process types:

1) process of a retinal image of such-and-such specific characteristics leading to a belief that there is a maple tree nearby
2) process of relying on a leaf shape to form a tree-classifying judgment
3) the visual process
4) vision in bright sunlight
5) perceptual process that occurs in middle-aged men on Wednesdays
6) process which results in justified beliefs[3]
7) perceptual process of classifying by species a tree behind a solid obstruction

There seems to be no end to the list of possible types for the token in question. To see that these types vary in reliability, consider type (7) and all of the process tokens that fall under it. In general, trying to classify trees according to species when those trees lie behind solid obstructions is not going to result in a high proportion of true beliefs. If this unreliable process type were to determine the justification of Smith's belief, his belief would be unjustified. In Smith's case, however, the solid obstruction is a glass window. Since the process type 'perceptual process of classifying a tree by species while looking through a glass window' would presumably be more reliable, Smith's belief would be justified if this were the relevant process type. Since Smith's belief seems to fall under both of these types—not to mention (1) through (6)—it is unclear whether his belief would be justified or unjustified according to reliabilism.

When Alvin Goldman (1992a) initially appealed to process types as the bearers of reliability, he claimed that process type selection must avoid two extremes. The first extreme he called the 'Single Case Problem,' which occurs when a process type is described so narrowly that only one instance of it ever occurs, and hence the type is either completely reliable or completely unreliable (Goldman 1992a, p. 115). If the result of the single instance is a true belief, then the ratio of true outputs to total outputs is 1, yielding complete reliability. A single false output can generate complete unreliability in the same manner. According to Goldman, process type selection must also avoid the 'No Distinction Problem,' which Feldman (1985, p. 161) characterizes as arising

> when beliefs of obviously different epistemic status are produced by tokens that are of the same (broad) relevant type. For example, if the relevant type for every case of inferring were the type "inferring," then [reliabilism] would have the unacceptable consequence that the conclusions of all inferences are equally well justified (or unjustified) because they are believed as a result of processes of the same relevant type.

Feldman relates the Single Case, No Distinction, and Generality Problems as follows.

> The problem for defenders of the reliability theory, then, is to provide an account of relevant types that is broad enough to avoid The Single Case Problem but not so broad as to encounter The No-Distinction Problem. Let us call the problem of finding such an account "The Problem of Generality." (op cit.)

Reliabilism thus requires:

i) for any process token, that there be a single process type that is the epistemologically relevant type under which that process token falls;
ii) that the relevant process type avoid both the Single Case and the No Distinction Problems; and
iii) that the reliability of the relevant process type be the primary determinant of the justification of all the beliefs produced by the process tokens that fall under that type.

The Generality Problem can be understood as the claim that conditions (i) through (iii) are never satisfied.

If process reliabilism cannot provide a plausible way to isolate relevant process types for particular cases of belief, it is, according to its critics, at best "radically incomplete" (Conee & Feldman 1998, p. 3) or at worst fraught with crippling conceptual difficulties. Conee and Feldman (op cit., pp. 5, 24) write,

> Our thesis is that the prospects for a solution to the generality problem for process reliabilism are worse than bleak.... Consequently, process reliability theories of justification and knowledge look hopeless.

In this paper I critically examine the Generality Problem and argue that it does not succeed as an objection to reliabilism. Although those who urge the Generality Problem are correct in claiming that any process token can be given indefinitely many descriptions that pick out indefinitely many process types, they are mistaken in thinking that reliabilists have no principled way to distinguish between relevant and irrelevant process types. My solution to the problem proceeds in two stages. In the first stage I present and discuss a set of necessary conditions that any process type must satisfy if it is to count as the relevant process type for some process token. I call this set the 'tri-level condition.'

> *The tri-level condition*
> The reliability of a cognitive process type *T* determines the justification of any belief token produced by a cognitive process token *t* that falls under *T* only if all of the members of *T*:
> a) solve the same type of information-processing problem *i* solved by *t*;
> b) use the same information-processing procedure or algorithm *t* used in solving *i*; and
> c) share the same cognitive architecture as *t*.

According to the tri-level condition, cognitive process types are information-processing types that are partially defined by their computational and algorithmic properties.

The tri-level condition significantly reduces the field of potentially relevant cognitive process types and, thus, goes some way toward solving the Generality Problem. However, the tri-level condition is not sufficiently strong to solve the Generality Problem by itself. Let *A* be the broadest process type that satisfies the tri-level condition with respect to some process token *t*. Using all manner of epistemologically relevant and irrelevant properties, *A* can be partitioned in indefinitely many ways such that *t* will fall within indefinitely many distinct subclasses of *A*. Since *A* satisfies the tri-level condition, any subclass of *A* will also satisfy the tri-level condition.

And since each subclass of *A* picks out a distinct type, *t* will fall under indefinitely many types that satisfy the tri-level condition. In other words, the Generality Problem will arise all over again. Nevertheless, tri-level condition does form an important part of an adequate solution to the problem by showing how process types are partially defined in terms of their computational and algorithmic properties and by ruling out type descriptions that do not pick out cognitive information-processing types at all.

The second stage of my solution to the Generality Problem answers the question of how to partition *A* into epistemologically relevant subclasses. Drawing upon Wesley Salmon's (1971, 1984) work on statistical explanation, I argue that the relevant process type for any process token *t* is the subclass of *A* which is the broadest objectively homogeneous subclass of *A* within which *t* falls. A subclass *S* is objectively homogeneous if there are no statistically relevant partitions of *S* that can be effected. The two stages of my solution combine to show that reliabilists can escape from the Generality Problem.

II. The Tri-Level Hypothesis

The tri-level condition is based upon the 'tri-level hypothesis,' which was first formulated by David Marr (1982) and has been explicated and developed by Zenon Pylyshyn (1984) and Michael Dawson (1998). The tri-level hypothesis has become the orthodox way for cognitive scientists to understand explanations of cognitive processing.[4] According to the tri-level hypothesis, there are three basic levels at which any cognitive behavior can and should be explained: the computational level, the algorithmic level, and the implementation level. An item of cognitive behavior cannot be adequately explained at only one (or even two) of these levels.

Some aspects of cognition are best captured at the computational level, where researchers ask questions like, "What information-processing problem is the system in question solving?" Explanations couched at the algorithmic level answer the question, "What method is the system using to solve this information-processing problem?" Implementation level theories deal with the question, "What physical properties are used to implement the (functional) method that the system uses to solve this information-processing problem?" (Dawson 1998, p. 288). These three sorts of explanation are pitched at different levels of abstraction, employ qualitatively different vocabularies, and capture different regularities. They are "tied together in an instantiation hierarchy, with each level instantiating the one above" (Pylyshyn 1984, p. 132). Dawson (1998, p. 33) claims,

> I would strongly argue that if you can omit one of the vocabularies without losing any predictive or explanatory power, then the system that you are explaining is not an information processor.[5]

Explanations at the implementation level have the advantage of being couched in the vocabulary of one of the "hard" sciences, such as physics or neurophysiology. Grounding psychological theories in physical explanations can give these theories a kind of legitimacy they might not otherwise have. Moreover, if one has a physical description of a system, one will be able to formulate predictions about that system's performance one would not likely be able to make with other sorts of description. This is particularly true if the predictions are based upon a knowledge of the limitations imposed by physical features of the system's components.

One limitation of implementation level explanations, however, concerns the multiple realizability of cognitive functions. Since a particular type of information-processing system can be given a variety of physical realizations, implementation level explanations will not be able to explain features that may be shared by differently realized systems. Another limitation of physical explanations of a system is that very few of a system's physical properties will be relevant to understanding the kind of information processing it is carrying out. In other words, only a small subset of a cognitive system's physically discriminable states will be computationally discriminable (cf. Pylyshyn 1984, p. 56). Since computational and algorithmic explanations are pitched at higher levels of abstraction, they are able to capture regularities and generalizations that implementation level theories cannot.

Algorithmic level theories provide us with (nonphysical) procedural explanations of the information-processing steps that are being carried out on particular systems and, as we shall see below, functional explanations of the fundamental cognitive architectures of those systems. This level of explanation abstracts from details about the kind of stuff out of which a machine is built. In fact, one can describe a machine at this level that has never been built at all. This level of explanation is important because explanations of what information-processing problem a system is solving (at the computational level) and what it is made out of (at the implementation level) may not tell us everything we want to know.[6]

Algorithmic explanations, however, are also subject to a limitation analogous to the one faced by implementation explanations: there are generalizations found at a more abstract level that they cannot capture. If two systems are solving the same information-processing problem by following different procedures, algorithmic explanations will fail to capture this commonality. Moreover, the purely formal nature of explanations at the procedural level do not typically tell us how to interpret the system's actions or the symbols that it is processing. An assignment or interpretation of the semantics of a system is going to be found at the computational level. Just as two systems can be solving the same information-processing problem by following different procedures, two systems can also be following the same procedure while solving different information-processing problems. For example, one system could be following a set of rules in the course of playing a game of chess while another system could be following the same set of rules (syntactically described) and yet be directing the course of an actual war. Because computational level explanations involve semantic interpretations, they are able to distinguish the two cases.

Computational level explanations, then, are indispensable in explaining cognition because they tell us what the states of a cognitive system represent and how to interpret the functional operations it is carrying out. Computational explanations, however, are also subject to limitations. As Dawson (1998, p. 94) claims, "Computational theories are purely semantic, in the sense that they account for what system states mean, but do not account for how these states come to be." They fail in this explanatory task in two ways: a) they cannot tell us how an output state results from the execution of a particular algorithm, given a certain input value; and b) they cannot tell us how the physical properties of the input and the internal state of the system result in the physical properties of the output state.

When constructing a scientific explanation of the behavior of a particular information-processing system, all three levels of explanation will be important. However, when one is interested (as we are) in features that are shared by all the members of a particular information-processing type, it becomes necessary to abstract away from many of the physical details of the system. The complete set of fundamental cognitive capacities common to all members of some information-processing type is known as that type of system's 'cognitive architecture.' Due to the multiple realizability of these capacities, descriptions of cognitive architectures are pitched at a functional level that ignores many of the physical details of their implementation.[7] The algorithmic level of explanation is usually divided into two parts to make room for descriptions of cognitive architectures: a) procedural explanations that account for the algorithm being executed by a cognitive system and b) cognitive architecture explanations that account for the kind of system executing the algorithm. Since we are interested only in features shared by every member of some process type, the tri-level condition does not accord any role to implementation level considerations. Although physical properties make important contributions to scientific explanations, they cannot help in selecting relevant cognitive process types.

III. The Tri-Level Condition

The tri-level condition requires that the process types whose reliability determines the justification of the belief tokens produced by the process tokens subsumed by those types be information-processing types. According to the tri-level condition, those information-processing types are defined by: a) the kinds of information-processing problems solved by token processes falling under those types; b) the information-processing procedures used to solve those problems; and c) the type of system executing those procedures.

8) Let I_1 be the property of solving information-processing problem i_1, I_2 the property of solving information-processing problem $i_2, \ldots,$ and I_s the property of solving information-processing problem i_s.
9) Let M_1 be the property of using procedure m_1 to solve an information-processing problem, M_2 the property of using procedure m_2 to solve an information-processing problem, ..., and M_t the property of using procedure m_t to solve an information-processing problem.
10) Let S_1 be the property of executing a problem-solving procedure on a system of type s_1, S_2 the property of executing a problem-solving procedure on a system of type $s_2, \ldots,$ and S_u the property of executing a problem-solving procedure on a system of type s_u.

Accordingly, each relevant process type T_i will be partially defined by a conjunction of the following properties.

i) Some I_j from the set of properties $I_1, \ldots, I_s$ $(1 \leq j \leq s)$
ii) Some M_k from the set of properties $M_1, \ldots, M_t$ $(1 \leq k \leq t)$
iii) Some S_h from the set of properties $S_1, \ldots, S_u$ $(1 \leq h \leq u)$

If a process type T_1 is partially defined by the conjunction of properties $I_1 \cdot M_1 \cdot S_1$, then T_1 can be the single relevant process type for two process tokens, t_1 and t_2, only if t_1 and t_2 both fall within I_1, M_1 and S_1.[8]

A. The tri-level condition can now be applied to the claims made by proponents of the Generality Problem. Recalling the story used in section I to illustrate the Generality Problem, Conee and Feldman offer the following description of the token process that produced Smith's maple tree belief.

> 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 s[t]imulation in combination with complex standing conditions in Smith's brain, result in Smith's forming the belief. This sequence of concrete events is the process [token] that caused the belief. (Conee & Feldman 1998, pp. 1–2)

There are two crucial features of this description. The first is that the token is described almost exclusively in physical terms. The second is that the only non-physical part of the description is the information that the token "result[s] in Smith's forming the belief." Conee and Feldman claim that the properties invoked in this description—which they take to be quite informative and detailed—are incapable of distinguishing any process type as the relevant one under which the process token falls. There is both an ounce of truth and a pound of confusion in this claim, and the tri-level condition can help sort things out.

One strand of the Generality Problem focuses on all of the physical properties that can be included in a description of a process token and claims—correctly—that these properties alone are incapable of picking out the relevant process type for the case in question. For any process type with physical realization *Z*, there will be indefinitely many algorithms that can be instantiated by *Z* and indefinitely many information-processing problems that can be solved by a process token of that physical type. In other words, the purely physical properties found in implementation level description are simply incapable of picking out relevant cognitive process types. The tri-level condition can explain why this insight does not lead to any serious problem for reliabilism. Process types are defined by computational and algorithmic properties rather than by physical properties because of the multiple realizability of any given process type. The inability of purely physical characterizations to capture cognitively relevant facts is precisely why cognitive scientists look to higher-order functional descriptions that abstract from many of the physical details of cognitive systems.

The appropriate conclusion to draw from the fact that the physical properties in Conee and Feldman's process token description are incapable of isolating a single relevant process type for the token that produced Smith's maple tree belief is that they have offered an inadequate process description. The severe limitations of purely physical descriptions of cognitive processes provide no support for the conclusion that reliabilism suffers from a debilitating defect.

Consider the second crucial feature of Conee and Feldman's description above, viz., that it includes the additional information that the token "result[s] in Smith's forming the belief." I agree that resulting in Smith's belief is a sufficient condition for belonging to a cognitive information-processing type. However, if this is the only thing one knows about the information-processing features of that token, one will not have a very clear idea about what information-processing problem the token process solved, and one will be completely ignorant about which algorithm—out of all the possible algorithms that could have been used to solve the problem—was actually used on the occasion in question. In other words, one will be totally in the dark about that process token's computational and algorithmic properties. It should be no wonder that anyone in this situation will be unable to pick out just one process type as the relevant one for the target process token.

In order for proponents of the Generality Problem to establish the conclusion that "the prospects for a solution to the generality problem for process reliabilism are worse than bleak" (Conee & Feldman 1998, p. 5), they must show more than that someone who is ignorant of the computational and algorithmic properties of a cognitive process token will be unable to determine the defining features of the single relevant process type for that token. They must show that there are no properties that make only one type relevant for each belief and process token pair.

As we have just seen, those who urge the Generality Problem often focus on the inability of a set of physical properties to determine relevant process types. These critics also look to the sort of properties found in computational level descriptions and make an analogous point. Critics display something like the list of process type descriptions used in section I to introduce the Generality Problem and claim—correctly—that there is nothing about the properties in these descriptions that seems to make one process type stand out as the relevant one for the target process token. This can be understood as the correct claim that computational level properties alone cannot determine relevant cognitive process types. A cognitive process type, however, is defined not only by its computational properties but also by its algorithmic properties—viz., the property of its members executing a particular type of information-processing procedure and the property of executing that procedure on a particular type of cognitive architecture. Instead of underwriting Conee and Feldman's (1998, p. 24) claim that the prospects for reliabilism are "hopeless," this line of reasoning serves merely to emphasize the complexity of the set of properties which define any relevant information-processing type.

We can see that the most common strategy used by Conee and Feldman to lodge the Generality Problem involves: i) focusing on the properties found at only one level of explanation (e.g., the implementation or the computational); and ii) arguing that none of the properties that figure in single-level explanations are capable of distinguishing one process type as the relevant type for each process token. The tri-level condition reveals why part (ii) of this strategy is based upon a genuine insight and why part (i) rests upon a mistake.

B. The tri-level condition can also help us sort through the list of candidate process types offered above for Smith's maple tree belief. We said that the process token that produced Smith's maple tree belief is a member of (at least) the following process types:

1) process of a retinal image of such-and-such specific characteristics leading to a belief that there is a maple tree nearby
2) process of relying on a leaf shape to form a tree-classifying judgment
3) the visual process
4) vision in bright sunlight
5) perceptual process that occurs in middle-aged men on Wednesdays
6) process which results in justified beliefs
7) perceptual process of classifying by species a tree behind a solid obstruction

How do these process type descriptions fare with respect to the tri-level condition? Only (1) and (2) give us much of an idea of the information-processing problem being solved by the token process. They both tell us that some kind of object-recognition is taking place on the basis of some kind of perceptual input. But these computational descriptions are much too vague to be very informative. Moreover, none of the process type descriptions in (1) through (7) give us an inkling as to how the target process went about solving the problem or what kind of system executed the information-processing procedure—i.e., none include algorithmic properties. In short, none of the process type descriptions provides us with the computational and algorithmic properties that partially define any cognitive process type—much less the single, epistemologically relevant process type we are seeking.

Description (5) poses a special challenge that the tri-level condition cannot solve. Its creator, Feldman, intentionally included obviously irrelevant properties to see whether reliabilists had the resources to exclude it from consideration. The tri-level condition claims that epistemologically relevant process types are partially defined by certain computational and algorithmic properties that are essential to cognitive information-processing. However, it does not require that *only* essential features should figure in the definition of a process type. The second stage of my solution to the Generality Problem—to which I now turn—is designed to deal with the challenge posed by cognitively irrelevant properties. In spite of the fact that the tri-level condition cannot exclude process type descriptions like (5), the tri-level condition still plays an important role in solving the Generality Problem by showing how cognitive process types are partially defined by their computational and algorithmic properties.

IV. Statistical Relevance

Let A be the broadest process type that satisfies the tri-level condition for some process token t. In addition to having the properties required by the tri-level condition (i.e., in addition to possessing the properties required for membership in A), t and the other members of A will have other properties as well. Some of these further properties will be relevant to the kind of information-processing carried out by members of A, while others will not. Variations in any of these properties will result in distinct types. If, for example, we let 'W' denote the property of occurring on Wednesdays and 'M' the property of occurring on Mondays, then some members of A will also be members of the type $A \cdot W$, and others will be members of $A \cdot M$. If A satisfies the tri-level condition, then $A \cdot W$ will. In fact, for any property F, $A \cdot F$ will satisfy the tri-level condition if A does. Consequently, for any process token, there will be indefinitely many process types that satisfy the tri-level condition and subsume the process token in question.

The question we must now face is, Why should some process type $A \cdot F$, rather than any other type $A \cdot G$, be the single relevant type for some cognitive process token t? If there is no way to make distinctions of relevance among all of the cognitive process types that satisfy the tri-level condition with respect to t, we will be left without a solution to the Generality

Problem. The second stage of my solution to the Generality Problem answers the challenge at hand. Drawing upon Wesley Salmon's (1971, 1984) work on statistical explanation, I argue that the relevant process type for some t is the subclass of A which is the *broadest objectively homogeneous subclass of A within which t falls*. A subclass S is objectively homogeneous if there are no statistically relevant partitions of S that can be effected.

To set up the second stage of my solution, I need to recast the Generality Problem in terms of conditional probability. To begin with, the reliability of some process type T_i is equal to the probability that a cognitive process token t_j will produce a true belief, given that t_j belongs to T_i.[9] Where '$R(x)$' means the reliability of x, we have

11) $R(T_i) = P(t_j \text{ produces a true belief} \mid t_j \text{ belongs to } T_i)$.

Where A is the broadest process type that satisfies the tri-level condition with respect to some process token t, the Generality Problem can be understood as the claim that:

a) A can be partitioned in indefinitely many ways such that t will fall into indefinitely many distinct subclasses of A;
b) for any subclass of A, $A \cdot F$, there will be another subclass of A, $A \cdot G$, such that
$P(t \text{ produces a true belief} \mid t \text{ belongs to } A \cdot F) \neq P(t \text{ produces a true belief} \mid t \text{ belongs to } A \cdot G)$; and
c) reliabilists have no principled way to determine which subclass of A represents the relevant type for the target process token.

In order to determine which subclass of A picks out the relevant process type for t, we need to provide a statistically relevant partition of A. That is, we must invoke a set of statistically relevant factors $C_1, \ldots, C_s$ that partitions A into a set of mutually exclusive and exhaustive cells $A \cdot C_1, \ldots, A \cdot C_s$. Statistical relevance can be understood as a comparison of probabilities (cf. Salmon 1971; 1984, p. 33). A condition C_j is a statistically relevant factor to the occurrence of B under circumstances A if and only if

12) $P(B|A) \neq P(B|A \cdot C_j)$

or, equivalently,

13) $P(B|A \cdot C_j) \neq P(B|A \cdot {—}C_j)$.

Each of the cells $A \cdot C_j$ in such a partition must be objectively homogeneous with respect to B; that is, none of the cells in the partition can be further subdivided in any manner relevant to the occurrence of B.

Let D be the property of occurring up close to a perceived object, O the property of occurring when the surface of a perceived object with the greatest area is oriented perpendicularly to the line between subject and object, L the property of occurring in a high degree of ambient lighting, E the property of occurring after less than one second of exposure to a perceived object, W the property of occurring on Wednesdays, M the property of occurring in middle-aged men, and B the property of producing a true belief. It is plausible to think that the following claims are true.

14) $P(B|A{\cdot}D) \neq P(B|A{\cdot}D{\cdot}O) \neq P(B|A{\cdot}D{\cdot}O{\cdot}L)$
15) $P(B|A{\cdot}D) \neq P(B|A{\cdot}D{\cdot}E)$
16) $P(B|A{\cdot}D{\cdot}O) \neq P(B|A{\cdot}D{\cdot}O{\cdot}E)$
17) $P(B|A{\cdot}D{\cdot}O{\cdot}L) \neq P(B|A{\cdot}D{\cdot}O{\cdot}L{\cdot}E)$
18) $P(B|A{\cdot}D{\cdot}O{\cdot}L) \neq P(B|A{\cdot}\,{-}D{\cdot}\,{-}O{\cdot}\,{-}L)$
19) $P(B|A{\cdot}D{\cdot}O{\cdot}L) = P(B|A{\cdot}D{\cdot}O{\cdot}L{\cdot}W{\cdot}M)$

In other words, it is plausible to suppose that factors D, O, L and E are statistically relevant but not that W and M are.

In order to provide a statistically relevant partition of A, the cells must not only be homogeneous; they must also be *maximally* homogeneous. A subclass S is maximally homogeneous if there is no larger class of which S is a subset and which is also objectively homogeneous. In other words, for some partition in terms of $F_1, \ldots, F_u$, $P(B|A{\cdot}C_j)$ cannot, in general, equal $P(B|A{\cdot}C_j{\cdot}F_h)$, where $1 \leq h \leq u$.[10] This requirement assures us that our partition $\{F_h\}$ does not introduce any irrelevant subdivision into the initial set of subclasses $A{\cdot}C_1, \ldots, A{\cdot}C_s$.

Each of the cells in the reference class partition must also be *objectively homogeneous*. The notion of objective homogeneity can best be illustrated by contrasting it with other varieties of homogeneity. According to Salmon (1966, p. 92), a class is *epistemically homogeneous* when we "may suspect that a given reference class is inhomogeneous, but not know of any way to make a relevant partition of it." This does "not demand that no possibility of a relevant partition can exist unbeknown to us" (Salmon 1984, p. 41). In cases where we know that a reference class is inhomogeneous but carrying out a relevant subdivision would be impractical, we can say that the reference class is *practically homogeneous*. We may decide not to carry out certain subdivisions because of the size of the class, the amount of statistical evidence available, the cost involved in getting more data, the difficulty of effecting the relevant subdivision, and how much is at stake in our decision (Salmon 1966, p. 92). By contrast, objective homogeneity is independent of both our knowledge and our practical interests. A reference class is objectively homogeneous if there is, in fact, no statistically relevant partition of it that can be made.

Reliabilists should reject any epistemically relativized version of homogeneity. According to externalist theories like reliabilism, the conditions for knowledge (or justified belief) are such that no one needs to know whether they are fulfilled in order for knowledge (or justified belief) to be possible. In other words, knowledge—like truth—is recognition-transcendent. One can know (or have the justified belief) that p without knowing that one knows that p (or having the justified belief that one's belief that p is justified). In the case of reliabilism, this means that agents do not have to know that their processes are reliable in order to have justified beliefs. Their processes must simply *be* reliable. Since reliabilists claim that reliability and epistemic justification are both recognition-transcendent properties, they cannot accept any account of these properties that does not preserve that recognition-transcendence. Consequently, only objective homogeneity will do.

If A is partitioned into a set of maximally and objectively homogeneous subclasses $A{\cdot}C_1, \ldots, A{\cdot}C_s$, the value of $P(t_j$ produces a true belief$|t_j$ belongs to $A{\cdot}C_j)$ will be the same for every t_j subsumed by each $A{\cdot}C_j$. Each cell in such a partition will be the relevant process type for every process token that falls under it.

V. Statistical Relevance and the No Distinction Problem

Recall for a moment the No Distinction Problem, which arises when beliefs of obviously different epistemic status are produced by process tokens grouped together into the same broad process type. Feldman (1985, p. 161) illustrates the No Distinction Problem with the following example.

> [I]f the relevant type for every case of inferring were the type 'inferring,' then [reliabilism] would have the unacceptable consequence that the conclusions of all inferences are equally well justified (or unjustified) because they are believed as a result of processes of the same relevant type.

I submit that the following two claims are identical.

20) No process type that *falls prey to the No Distinction Problem* can be the process type whose reliability determines the justification of beliefs produced by process tokens of that type.
21) No process type that *can be partitioned in statistically relevant ways* can be the process type whose reliability determines the justification of beliefs produced by process tokens of that type.

In other words, the injunction to avoid the No Distinction Problem is nothing more than the injunction to avoid statistically inhomogeneous process types. The process type 'inference' is obviously not homogeneous because it includes not only *modus ponens* and *modus tollens* but also affirming the consequent and denying the antecedent. There are statistically relevant partitions of that type that can and should be made. The same thing can be said for another of Feldman's examples. Feldman (1993, p. 41) suggests that any token visual perceptual process will belong to each of the following process types.

22) the visual process
23) the perceptual process
24) the cognitive process

Feldman notes that each process type clearly falls victim to the No Distinction Problem, since each "process" is actually a family of process types. Feldman does not, however, see that reliabilists can take the further step suggested by his remarks—viz., divide each class into maximally homogeneous subclasses until each subclass no longer faces the No Distinction Problem.

Feldman and Goldman have both claimed that an adequate solution to the Generality Problem must avoid both the Single Case Problem and the No Distinction Problem. What Feldman and Goldman fail to realize is that the formulation of the No Distinction Problem contains within it resources reliabilists need to solve the Generality Problem. If we think of the Single Case and No Distinction Problems as "extremes," we may be misled into thinking that avoiding the two extremes will not be sufficient for isolating unique types. Within the area that falls between the two extremes, it may seem as if there will be an abundance of process types for any given process token and that the Generality Problem will continue to pose a serious threat to reliabilism. On the contrary, I claim that repeated applications of the maxim 'Avoid the No Distinction Problem' can partition the broadest process type satisfying the tri-level condition within which some process token falls into (maximally and objectively) homogeneous subclasses, where the subclass which subsumes the process token in question will be the relevant process type for that token. In short, the two stages of my proposal combine to solve the Generality Problem.[11]

VI. Conclusion

Proponents of the Generality Problem claim that each process token belongs to indefinitely many process types and that reliabilists have no principled way to isolate a single relevant process type for each process token. According to reliabilism, the justification of a belief is a direct function of the reliability of the relevant process type under which the token process that produced that belief falls. The greater the reliability of that process type, the greater the justification of the target belief. If there is not a single relevant process type for each process token, then there will be no fact of the matter regarding the justification of the beliefs produced by those tokens.

My solution to the Generality Problem has proceeded in two stages. The first stage presents a set of necessary conditions that partially define relevant cognitive process types. According to the tri-level condition, each relevant process type is partially defined by the type of information-processing problem solved by its members, the type of procedure used to solve that problem, and the particular kind of system used to execute that procedure. These computational and algorithmic properties ensure that the only candidates for relevant process types will be information-processing types, and they tell us what broad information-processing type a given process belongs to.

We have seen that the broadest process type *A* that satisfies the tri-level condition with respect to some process token *t* can be partitioned in indefinitely many ways such that *t* can be made to fall within indefinitely many distinct subclasses of *A*. In the second stage of my solution to the Generality Problem I have argued that, for any process token *t* falling within *A*, the single relevant process type for *t* is the broadest objectively homogeneous subclass of *A* that subsumes *t*.

My solution to the Generality Problem can be summed up very simply. If we are dealing with cognitive information-processing types (partially defined by their computational and algorithmic properties), then we should let the maxim 'Avoid the No Distinction Problem' be our guide. While one might think that we should first determine which process types are the relevant ones and then determine the reliability of those types, my solution suggests that we let differences in degree of reliability (and unreliability) be our guide to relevance. The two stages of my solution combine to show that reliabilists have a principled way to distinguish between relevant and irrelevant process types.[12]

Notes

[1] In the exposition and defense of reliabilism that follows, I take Alvin Goldman's (1986, 1992a, 1994) account of justified belief to be the most representative and developed statement of process reliabilism.

[2] The objection can be found in the following surveys of contemporary epistemology: Crumley (1999, pp. 76–79), Hetherington (1996, pp. 40–41), Lycan (1988, pp. 110–111), Plantinga (1993, p. 198), Pollock (1986, pp. 118–120), and Pollock and Cruz (1999, pp. 116–8). For further discussion of the Generality Problem, see Heller (1995), Schmitt (1992, ch. 6), Alston (1995), Wallis (1994), Baergen (1995, ch. 4), and Brandom (1994, ch. 4). Goldman (1992b)—the foremost defender of reliabilism—claims in a recent reference work devoted to epistemology that the Generality Problem is one of the most pressing challenges facing reliabilism.

[3] (6) employs an evaluative term, 'justified.' According to reliabilism, a belief can be justified only if the single relevant process type that the token process falls under is reliable. The determination of relevant types, then, is prior to determinations of justification, so (6) cannot adequately serve to mark out a relevant type.

[4] Connectionists, of course, would challenge this claim to orthodoxy. For a fascinating defense of how connectionists cannot avoid the claims of the tri-level hypothesis, cf. Dawson (1998, ch. 5).

[5] Connectionists would not deny that complete explanations of their systems should include both computational and implementation level descriptions. They would, however, object to the tri-level hypothesis' claim that algorithmic level explanations are required. Connectionists systems, they claim, do not rely on algorithms to solve information-processing problems. For a compelling challenge to this connectionist claim, cf. Dawson (1998, pp. 134–142).

[6] Consider, for example, Deep Blue, which in 1995 became the first chess-playing computer to beat a reigning world chess champion under tournament conditions. Deep Blue consists of four specialized computers processing in parallel, enabling it to consider one billion moves per second. At the computational level of explanation, Deep Blue and Gary Casparov can be described as solving the same information-processing problem, viz., playing world championship chess. They will obviously have different physical or implementation level explanations, since one is mostly silicon and plastic and the other is flesh and bones. But the differences between them are more than merely physical. They also do not follow the same algorithms or procedures. Although no one currently has a complete understanding of the nature of the information-processing in a human chess champion like Casparov, we can be quite certain that he isn't going through the same information-processing steps as Deep Blue. For further discussion cf. Dawson (1998).

[7] Brain researchers agree that there are simply too many connections between neurons and too few genes in the human genome for precise information regarding all of the neural pathways to be stored in our DNA. Even if our DNA contained information about nothing else than connections between neurons, there would still not be enough storage space in the human genome for all of the relevant information. Identical twins, who share the same genetic code, do not even have the physically same neural pathways.

[8] Another way to look at the tri-level condition is this: the tri-level condition claims that a process type T is the relevant type for two process tokens only if: i) those two process tokens are 'strongly equivalent'; and ii) those tokens are strongly equivalent by virtue of falling within T. According to Fodor (1968), Pylyshyn (1984) and Dawson (1998, pp. 98, 173), two systems that are solving the same information-processing problem using possibly different methods are 'weakly equivalent.' 'Strongly equivalent' systems employ identical procedures to solve the same information-processing problem. The latter are not merely input-output equivalent but are identical systems running the same program, as it were. For discussion of how to determine whether two cognitive systems are strongly equivalent, cf. Pylyshyn (1981; 1984, ch. 5; 1989), Dawson (1998, p. 110ff.), Cummins (1983, ch. 1), Sternberg (1995), and Newell and Simon (1972).

[9] Note that we are not concerned with the question, "What is the prior probability that a certain belief token is true?" If a process type is 90% reliable, that does not mean that each of the beliefs produced by the token processes of that type will have a .9 prior probability of being true. Those beliefs may vary widely in their independent probabilities. Some, in fact, may have a probability value of less than .5. Reliability is concerned only with the likelihood that a belief will be true, given that it is produced by a process token that belongs to a certain belief-forming process type.

[10] An exceptional but permissible case mentioned by Salmon concerns a coin toss experiment with one fair and two biased coins. The probability of heads with one of the biased coins C_1 is .9 and the probability of tails with the other biased coin C_2 is .9. The probability of heads with the fair coin C_3 is .5. The probability of heads, given that one randomly selects one of the coins, will also be .5. In other words, the probability of heads, given the class $\{C_1, C_2, C_3\}$, is equal to the probability of heads, given $\{C_3\}$. This equality is permissible, since partitioning the class $\{C_1, C_2, C_3\}$ into three subsets would be statistically relevant and would result in an overall increase of homogeneity of the subclasses.

[11] The second stage of my solution to the Generality Problem follows Salmon's (1971, 1984) solution to the 'reference class problem' that arises in statistical explanation. The reference class problem is analogous to the Generality Problem and is even characterized in terms of avoiding two extremes that strongly resemble the Single Case and No Distinction Problems. The reference class

problem arises for the frequency interpretation of probability when we try to ascertain the probability of a single event because there are indefinitely many ways to assign some x, which falls in the attribute class, to a reference class. Salmon (1971, pp. 41–42) writes,

> The reference class must, therefore, be broad enough to provide the required number of instances for examination to constitute evidence for an inductive inference [from observed frequency to relative frequency]. At the same time, we want to avoid choosing a reference class so broad that it includes cases irrelevant to the ones with which we are concerned.

The first pitfall Salmon mentions is analogous to the Single Case Problem and the second pitfall is analogous to the No Distinction Problem.

Interestingly enough, some of the solutions to Generality Problem that have been proposed strongly resemble some of the solutions offered to the reference class problem. Hempel (1965) suggested that the reference classes that figure in the statistical laws of statistical explanations be maximally specific. Hempel's 'requirement of maximal specificity' (RMS) demands that

> when the class to which the individual case is referred for explanatory purposes . . . is chosen, we must not know how to divide it into subsets in which the probability of the fact to be explained differs from its probability in the entire class. (Salmon 1984, p. 29)

In similar fashion, William Alston (1995, p. 14) tries to solve the Generality Problem by suggesting that appropriate process types are those that are 'maximally specific.' I would be surprised if Alston's own maximal specificity requirement was not inspired in some way by Hempel's. The suggestions of Hempel and Alston are both incapable of keeping irrelevant properties out of reference class or process type descriptions.

[12] This paper has benefited greatly from comments given by James Bohman, Eleonore Stump, Chris Pliatska, Husain Sarkar, and two anonymous reviewers from *Noûs*.

References

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

Baergen, R. (1995) *Contemporary Epistemology* (Ft. Worth, TX: Harcourt Brace).

Brandom, R. B. (1994) *Making it Explicit* (Cambridge, MA: Harvard University Press).

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

Crumley, J. S. (1999) *An Introduction to Epistemology* (Mountain View, CO: Mayfield).

Cummins, R. (1983) *The Nature of Psychological Explanation* (Cambridge, MA: MIT Press).

Dawson, M. R. W. (1998) *Understanding Cognitive Science* (Oxford: Blackwell).

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

Feldman, R. (1993) "Proper Functionalism," *Noûs* 27: 34–50.

Fodor, J. A. (1968) *Psychological Explanation* (New York: Random House).

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

Goldman, A. I. (1992a) "What is Justified Belief?" in *Liaisons: Philosophy Meets the Cognitive and Social Sciences* (Cambridge, MA: MIT Press).

Goldman, A. I. (1992b) "Reliabilism," in J. Dancy & E. Sosa (eds.) *A Companion to Epistemology* (Cambridge, MA: Blackwell).

Goldman, A. I. (1994) "Naturalistic Epistemology and Reliabilism," in P. French, T. Uehling, T. Wettstein (eds.) *Midwest Studies in Philosophy XIX: Philosophical Naturalism* (Minneapolis: University of Minnesota Press).

Heller, M. (1995) "The Simple Solution to the Problem of Generality," *Noûs* 29: 501–515.
Hempel, C. G. (1965) "Aspects of Scientific Explanation," in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science* (New York: The Free Press).
Hetherington, S. C. (1996) *Knowledge Puzzles: An Introduction to Epistemology* (Boulder, CO: Westview).
Lycan, W. G. (1988) *Judgement and Justification* (New York: Cambridge University Press).
Marr, D. (1982) *Vision* (San Francisco: W. H. Freeman).
Newell, A. & Simon, H. A. (1972) *Human Problem Solving* (Englewood Cliffs: NJ: Prentice-Hall).
Plantinga, A. (1993) *Warrant: The Current Debate* (Oxford: Oxford University Press).
Pollock, J. L. (1986) *Contemporary Theories of Knowledge* (Totowa, NJ: Rowman and Littlefield).
Pollock, J. L. & Cruz, J. (1999) *Contemporary Theories of Knowledge*, 2nd edition (Lanham, MD: Roman and Littlefield).
Pylyshyn, Z. (1981) "Complexity and the Study of Artificial and Human Intelligence," in J. Haugeland (ed.) *Mind Design* (Cambridge, MA: MIT Press).
Pylyshyn, Z. (1984) *Computation and Cognition: Toward a Foundation for Cognitive Science* (Cambridge, MA: MIT Press).
Pylyshyn, Z. (1989) "Computing in Cognitive Science," in M. I. Posner (ed.) *Foundations of Cognitive Science* (Cambridge, MA: MIT Press).
Salmon, W. C. (1966) *The Foundations of Scientific Inference* (Pittsburgh: University of Pittsburgh Press).
Salmon, W. C. (1971) "Statistical Explanation," in W. C. Salmon, R. Jeffrey, & J. Greeno *Statistical Explanation and Statistical Relevance* (Pittsburgh: University of Pittsburgh Press).
Salmon, W. C. (1984) *Scientific Explanation and the Causal Structure of the World* (Princeton, NJ: Princeton University Press).
Schmitt, F. (1992) *Knowledge and Belief* (New York: Routledge, Chapman and Hall).
Sternberg, S. (1995) "Discovering Mental Processing Stages: The Method of Additive Factors," in D. Scarborough & S. Sternberg (eds.) *An Invitation to Cognitive Science*, 2nd edition, vol. 4 *Methods, Models, and Conceptual Issues* (Cambridge, MA: MIT Press).
Wallis, C. (1994) "Truth-Ratios, Process, Task, and Knowledge," *Synthese* 98: 243–269.