



LUND UNIVERSITY

Is There a Statistical Solution to the Generality Problem?

Dutant, Julien; Olsson, Erik J

Published in:
Erkenntnis

DOI:
[10.1007/s10670-012-9427-y](https://doi.org/10.1007/s10670-012-9427-y)

2013

[Link to publication](#)

Citation for published version (APA):

Dutant, J., & Olsson, E. J. (2013). Is There a Statistical Solution to the Generality Problem? *Erkenntnis*, 78(6), 1347-1365. <https://doi.org/10.1007/s10670-012-9427-y>

Total number of authors:

2

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Is there a Statistical Solution to the Generality Problem?

Julien Dutant and Erik J. Olsson

Abstract: This article is concerned with a statistical proposal due to James R. Beebe for how to solve the generality problem for process reliabilism. The proposal is highlighted by Alvin I. Goldman as an interesting candidate solution. However, Goldman raises the worry that the proposal may not always yield a determinate result. We address this worry by proving a dilemma: either the statistical approach does not yield a determinate result or it leads to trivialization, i.e. reliability collapses into truth (and anti-reliability into falsehood). Various strategies for avoiding this predicament are considered, including revising the statistical rule or restricting its application to natural kinds. All amendments are seen to have serious problems of their own. We conclude that reliabilists need to look elsewhere for a convincing solution to the generality problem.

1. Introduction

Process reliabilism, or reliabilism for short, is the view that S knows that p if and only if (i) p is true, (ii) S believes that p , (iii) S 's belief that p was acquired through a reliable process, and (iv) an appropriate anti-Gettier condition is satisfied. Reliabilism is sometimes advocated as a theory of epistemic justification, the main idea being that a person is justified in belief that p just in case her belief that p was formed via a reliable process, i.e., the process by means of which the belief was acquired was reliable.¹ For the

¹ The process reliabilist account of knowledge was originally formulated by Ramsey (1931). For a modern (post-Gettier) account, see Goldman (1986). The process reliabilist theory of justification was first put forward in Goldman (1979).

purposes of the following discussion, there is no need to make a sharp distinction between these two brands of reliabilism.

The source of the longstanding *generality problem* for reliabilism is the observation that, because a process token is an unrepeatable causal sequence occurring at a particular time and place, it makes no good sense to ask whether a token process is reliable in itself. Rather, what can fundamentally be reliable or not are process *types*. A process token can still be said to be (un)reliable in a derivative sense if its associated process type is (un)reliable. For instance, the concrete process of Jones's coming to believe that he won the lottery on May 1, 2007, is itself neither reliable nor unreliable. However, given that its associated type is taken to be "belief formed through reading the local newspaper", it is (probably) reliable.

The generality problem now arises because each token process can be classified as belonging to a great many different types, and it is not obvious how to single out one of these types rather than another as *the* unique associated type of the process in question. For example, the process leading up to Jones's belief could be classified narrowly as belonging to the type whose sole member is Jones's coming to have his belief about the lottery, or, to take the other extreme, broadly as a belief formed through reading.

Furthermore, depending on what type is singled out as special, we may get different verdicts as regards the reliability of the process in question. Given the narrow classification in terms of the type whose sole member is the process producing Jones's belief, that process will, if the belief is true, be reliable. If it is seen instead as instantiating the general type "reading", it might be judged unreliable. Reading in general, irrespective of what is being read, is probably not reliable to an extent that suffices for knowledge.

These considerations reveal what appears to be an unacceptable lacuna in the reliabilist account of knowledge. From the reliabilist's perspective, whether a person knows or not will in many cases depend on whether the type of process producing the belief in question is reliable. And yet, as most commentators would agree, reliabilists have generally failed to clearly identify *the* type pertaining to a given token. In the absence of a principled account for how to select the relevant type, the reliabilist theory appears to be, in Conee and Feldman's words, "radically incomplete" (1998, p. 3).

Our concern in this paper will be with a recent proposal made by James R. Beebe (2004). The theory belongs to the tradition of strategies that invoke scientific types, as explained in the next section, but it differs from many other approaches in addressing the generality problem head on in a manner which is both principled and true to the spirit of reliabilism and externalism generally. Beebe's proposal has recently been highlighted by Alvin Goldman, the most distinguished current advocate of a reliabilist theory, as one of four candidate solutions deserving special attention (Goldman, 2008).² Goldman also notes a potential technical problem. We will return to Goldman's observation in a moment. Neither Beebe nor – to our knowledge – anyone else has investigated the precise consequences of the proposal. Before we take on this challenge, we will review some proposed desiderata for reasonable solutions to the generality problem as well as some candidate solutions, including a few quite recent ones.

2. Background

Reliabilists have not been insensitive to the generality problem which was identified in Goldman (1979) and is portrayed as a serious issue for reliabilism in Goldman (1986)

² See also Adler (2005) and Baumann (2011) for discussions of Beebe's proposed solution.

and, in particular, Goldman (2009). It is now considered to be a main challenge for a reliabilist theory by reliabilists and their critics alike.³ The problem has been stressed by Feldman (1985) and Conee and Feldman (1998). According to Feldman, solving the generality problem for reliabilism requires showing how to avoid the *single case problem* and the *no distinction problem*. The single case problem 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 depending on whether the belief is true or false. The no distinction problem arises when “beliefs of obviously different epistemic status are produced by tokens that are of the same (broad) relevant type” (Feldman 1985, p. 161; Beebe, p.179). These two problems were illustrated above in our example with Jones’s belief about the lottery.⁴

Conee and Feldman lay down three additional requirements for a solution to the generality problem. First, a solution must be “principled” in the sense of not being made on an *ad hoc* basis. Second, the rule must make reasonable epistemic classifications, by which is meant that the types identified must have a reliability that is plausibly correlated with the justificational status of the beliefs in question. Third, a solution must remain true

³ The generality objection can be found in many surveys of contemporary epistemology, e.g. Hetherington (1996), pp. 40-41, Lycan (1998), pp. 110-111, Plantinga (1993), p. 198, Pollock (1986), pp. 118-120, and Pollock and Cruz (1999), pp. 116-118. For a more recent example, see Lemos (2007), pp. 92-94.

⁴ The generality problem is usually stated as one of finding a *unique* relevant process type for each process token. Logically speaking, however, it would suffice to identify, for each token, a *class* of types that are either all reliable or all unreliable. With this caveat in mind, we will follow the mainstream and assume that the problem is to find a unique type for each token.

to the spirit of the reliabilist approach and not characterize the relevant type of process in epistemic terms that are alien to reliabilist theorizing.

Conee and Feldman also provide a helpful classification of different solutions to the generality problem in terms of *common sense*, *scientific* or *contextually determined* types. As for the first approach, it is tempting to classify belief-forming process in terms of categories like “confused reasoning”, “wishful thinking”, or “hasty generalization”. But, as Conee and Feldman note, reference to such commonsense types does not by itself solve the generality problem, for there are usually several common sense types that would fit a given token. Jones can be described as “reading a newspaper” or simply as “reading” where both descriptions plausibly pick out common sense types.

A second option is to identify an appropriate process type in scientific terms, i.e., by reference to “natural kinds”. A proposal along these lines was made by Alston (1995), who wrote:

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 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” (p. 11).

Yet, there is, Conee and Feldman insist, a uniqueness problem facing this approach as well (pp. 10-11):

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.

In his attempt to avoid this sort of criticism, Alston suggests narrowing down the different candidate types by appealing to the further criteria of *psychological realism* and *maximum specificity*. The idea is that the relevant type for a given process token corresponds to the maximally specific psychologically realistic natural kind to which the token belongs. In response, Conee and Feldman remark that “process reliability theories are supposed to appeal to much broader types” (1998, p. 16).

Mark Heller (1995) offers a contextualist approach to the generality problem. Heller believes that there is a sense in which the generality problem, as it has been commonly understood, is unsolvable. For there is, Heller maintains, no fixed principle for selecting the relevant level of generality. In so far as the reliabilist is urged to supply such a fixed principle, the demand is unreasonable. Rather, what counts as correct varies from context to context. Conee and Feldman (1998) as well as Goldman (2008) find it doubtful whether an appeal to contextual factors suffices to narrow down a unique type for every context.

Having surveyed the different proposals that figure in the literature, only to find that they all fail to comply with at least one desideratum, Conee and Feldman conclude that “the prospects for a solution to the generality problem for process reliabilism are worse than bleak” (p. 5) and that “[i]n the absence of a brand new idea about relevant types, the

problem looks insoluble” (p. 24), so that “process reliability theories of justification and knowledge look hopeless” (ibid.).

Since the appearance of Conee and Feldman’s critical survey a number of proposals have been put forward addressing the generality problem. Brandom (1998) advances a theory that is contextualist in the sense that it focuses on third-person attribution of reliability. However, it is doubtful whether Brandom thereby really solves the problem or mainly reformulates it in inferentialist terms. Wunderlich (2003) offers a technically highly sophisticated solution to the generality problem in which the justificatory status of a belief is not a function of a *single* appropriate type for each token but of a reliability *vector* associated with each token corresponding to the reliability values for the different possible types. Adler and Levin (2002) maintain that the generality problem is too general: if sound, it would apply also to processes mentioned in scientific explanations, which in their mind only shows that the generality problem is “illusory” to being with.⁵ Comesaña (2006) argues, similarly, that the generality problem is no problem unique to reliabilist but that it is one that many justification-based theories share, Conee and Feldman’s own evidentialist approach being a case in point. Bishop (2010) makes a related point, claiming that the generality problem is “everybody’s problem”. Kappel (2006) takes a stance that bears some resemblance to Heller’s, arguing that there are no facts that determine a relevant type, but that this is not a serious objection to the reliabilist theory.⁶

Although much work has gone into solving the generality problem, none of the proposals that have emerged so far has gained widespread acceptance (e.g. Bishop 2010,

⁵ In their response to Adler and Levin, Feldman and Conee (2002) reiterate their view that the generality problem is real and remains unsolved.

⁶ See also the exchange between Leplin (2007) and Christensen (2007).

p. 285). It could be added that a fair number of the candidate solutions do not really address the problem head on: they do not really show how token processes should be properly classified but focus rather on undermining the claim that the generality problem is a real and/or distinctive problem for reliabilism. As we will see, Beebe's statistical solution is interestingly different in this respect.

3. The statistical approach

Suppose that Smith looks out his window, sees a maple tree, and forms the belief that there is a maple tree nearby.⁷ The token process responsible for Smith's belief is an instance of (at least) the following class of 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 in forming 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 can we select a unique process type whose reliability is relevant for the justification of Smith's belief?

⁷ The example is due to Conee and Feldman (1998).

The first step in Beebe's solution is a "tri-level condition" that types must satisfy in order to be relevant. The condition draws on Marr's well-known tri-level hypothesis according to which any information processing system must be explained at three levels: computational, algorithmic and implementational (Marr, 1982). Beebe's idea is that a process type is relevant only if it is a uniform information-processing type (Beebe, p.183). This requires uniformity at the computational level: all tokens of the type should solve the same information-processing problem. It also requires uniformity at the algorithmic level: all tokens of the type should use the same information-processing algorithm, and the algorithm should be executed by systems sharing the same "cognitive architecture", that is, systems sharing the same "complete set of fundamental cognitive capacities" (ibid.). Uniformity at the implementational level, by contrast, is not required. Summing up, the tri-level condition states that "[t]he 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 ." (Beebe, p. 180)

Returning to the example with Smith, Beebe now claims that "none of the process type descriptions in 1. – 7. provide 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" (Beebe, p. 187). Types 1 and 2 come closest, as they at least hint that the problem solved is one of pattern recognition. Still, many distinct information-processing procedures could be applied to the problem. So the tokens of

these types are not algorithmically uniform. They are, at best, incomplete types that need to be supplemented with details about the relevant algorithm and cognitive architecture.

Let us grant, for the sake of the argument, that some contrived process types can be filtered out by appeal to the tri-level condition. Consider however the process type description 5.: “perceptual process that occurs in middle-aged men on Wednesdays”. Because this type contains what seem to be inessential features it would presumably not qualify as admissible in an intuitive sense. Can such features be filtered out through an appeal to the tri-level condition? Beebe does not think so. In his view, “for any property F , the property $A \& F$ will satisfy the tri-level condition if A does” (p. 187, notation adapted). Suppose A is a process type that satisfies the tri-level condition. By definition, all its instances satisfy conditions a), b), c). Now consider an irrelevant feature F such as “occurring in middle-aged men” or “occurring on Wednesdays”. Since all instances of the subtype $A \& F$ are also instances of A , they satisfy conditions a), b), and c). So the subtype $A \& F$ satisfies the tri-level condition. Hence if some perceptual process type A (suitably supplemented with algorithmic and other details) satisfies the tri-level condition, so will the conjoined type “instance of A that occurs in middle-aged men on Wednesdays”. Because there are often many ways in which a given admissible type can be made more specific by adding intuitively irrelevant features, Beebe concludes that “for any process token, there will be indefinitely many process types that satisfy the tri-level condition and subsume the process token in question” (Beebe, 187). Thus there is a need for a second filter that excludes types that have inessential parts.

The second step of Beebe’s proposal aims at screening out inessential features and selecting a single relevant type. It relies on the concept of statistical relevance: a condition C is a *statistically relevant* factor to the occurrence of B under circumstances A if and only if $P(B | A \& C) \neq P(B | A)$. Beebe goes on to define the reliability of a process

type T as the probability that a token process produces a true belief given that it is of that type: $P(t \text{ produces a true belief} \mid t \text{ belongs to } T)$. Now suppose that F is a feature that some instances of process T have. F is a statistically relevant feature for production of true belief by tokens of T iff the reliability of the subtype $T\&F$ differs from that of T , that is: iff $P(t \text{ produces a true belief} \mid t \text{ belongs to } T\&F) \neq P(t \text{ produces a true belief} \mid t \text{ belongs to } T)$. A type T is *homogenous* if “there are no relevant statistically relevant partitions of [T] that can be effected” (Beebe, p.181).⁸ This means that there is no subtype $T\&F$ of T such that F is statistically relevant to whether instances of T produce a true belief. Beebe’s final proposal is that the type relevant for evaluating a belief produced by a given token process t , among the types of t satisfying the tri-level condition, is the *broadest homogeneous type (with respect to the production of true beliefs) within which t falls* (Beebe, p.188).

The proposal is meant to work as follows. Suppose T is some visual process type that satisfies the tri-level condition. Suppose that its subtypes include “instances of T occurring in a high degree of ambient light” ($T\&L$), “instances of T occurring up close to a perceive object” ($T\&D$), as well as $T\&\sim L$, $T\&\sim D$, $T\&D\&L$, $T\&D\&\sim L$, $T\&\sim D\&L$ and $T\&\sim D\&\sim L$, and suppose that t is an instance of $T\&D\&L$. Presumably, D is statistically relevant to whether instances of T produce true beliefs (Beebe, p. 189): $P(t \text{ produces a true belief} \mid t \text{ belongs to } T\&D) \neq P(t \text{ produces a true belief} \mid t \text{ belongs to } T)$. So T is not a homogenous type. On Beebe’s view, T is therefore not the relevant type when evaluating the belief produced by t . We must move to a more specific type, such as $T\&D$. But again, L is presumably statistically relevant to whether instances of $T\&D$ produce true beliefs:

⁸ Beebe specifies that the homogeneity he has in mind is *objective* rather than *epistemic* (there are not *known* statistically relevant partitions) or *practical* (there are no *feasible, tractable* statistically relevant partitions.)

$P(t \text{ produces a true belief} \mid t \text{ belongs to } T\&D\&L) \neq P(t \text{ produces a true belief} \mid t \text{ belongs to } T\&D)$. So $T\&D$ is not a homogeneous type either. We must move to a still more specific type, such as $T\&D\&L$. However, we should eventually reach some stage at which further divisions are not statistically relevant. For instance, a strict subtype of $T\&D\&L$ “instances of $T\&D\&L$ occurring in Wednesday” ($T\&D\&L\&W$). Presumably, W is not statistically relevant to whether instances of $T\&D\&L$ produces true belief. Hence, Beebe claims, $T\&D\&L\&W$ will not be the *broadest* homogeneous type. So, Beebe’s proposal seems to have the desired effect of excluding types that have irrelevant parts.

Beebe thus summarizes his proposal as follows: “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.” (p.191) By “applying the maxim ‘Avoid the No Distinction problem’”, Beebe means to say: if a type can be partitioned in subtypes that are statistically relevant to the production of true belief, then we should move on to a more specific subtype. Beebe thinks that repeated applications of that process lead to a unique type, and that this unique type is the relevant one when assessing whether a belief is justified.

Still, Beebe makes no attempt in his article to prove that his proposal works. First, he does not show that his procedure will always select *some* type. This *existence worry* was raised by Alvin I. Goldman in a recent overview of reliabilism, in which he observes that “there remains the lingering question of whether there is always a set of conditions that meet Beebe’s standards, i.e., that generate an appropriate partition” (Goldman, 2008). Second, Beebe also fails to show that the procedure will select a *unique* type. Let us call this the *uniqueness worry*. Below we will address the existence worry by proving that *if*

the statistical approach yields a definite type, *then* reliability collapses into truth (and anti-reliability into falsity). We save the uniqueness worry for section 8.

4. A trivialization result

We will now show that even if Beebe's strategy solves the existence problem, then it leads to the collapse of reliability, and hence of justification, into truth. Since not all true beliefs are justified, the proposal fails to make epistemically reasonable classifications.

We will establish the following:

Observation 1: For any token process t producing a true belief, if T is a type of t selected by Beebe's constraints, then T is reliable.

Observation 2: For any token process t producing a false belief, if T is a type of t selected by Beebe's constraints, then T is not reliable.

Let t be some token belief-forming process and suppose T is a type of t that satisfies Beebe's constraints. Hence, T satisfies the tri-level condition and T is a broadest homogeneous type of t with respect to the production of true beliefs. For the first observation, suppose that t is true. Consider T_+ , the set of all instances of T that produced a true belief. Trivially, the subtype T_+ is perfectly reliable: $P(t \text{ produced a true belief} \mid t \text{ belongs to } T_+) = 1$. Now if T is not perfectly reliable, then T is not homogeneous: it has a subtype T_+ such that $P(t \text{ produced a true belief} \mid t \text{ belongs to } T_+) \neq P(t \text{ produced a true belief} \mid t \text{ belongs to } T)$. But we assumed that T is homogeneous. So T is perfectly reliable.⁹

⁹ In the simplest case, T is T_+ itself. But T may also be a type with an infinite number of true-belief instances and a finite number of false-belief instances, so that the probability that a belief is true given that it is an instance of that type is still 1.

For the second observation, suppose that t produced a false belief. Consider T^- , the set of all instances of T that produced a false belief. Trivially, T^- is perfectly anti-reliable. If T is not perfectly anti-reliable, it is not homogeneous. But we assumed T was homogeneous, so it must be perfectly anti-reliable. The conclusion is that if Beebe's strategy selects a type for token t , that type will be perfectly reliable if the belief produced by t is true, and it will be perfectly anti-reliable if the belief produced by t is false.

The upshot is that either Beebe's strategy fails to select a type, or it collapses reliability into truth. As a result, Beebe's strategy does not solve the Generality Problem. For there are unjustified true beliefs, e.g. true beliefs arising from wishful thinking. Either Beebe's constraints fail to select a type for those beliefs, in which case the strategy fails with regard to Goldman's existence worry, or they classify those beliefs as reliable and hence justified, in which case the strategy is epistemically unreasonable. As for the latter horn of the dilemma, we recall Conee and Felman's additional desideratum that a rule for typing token processes must make reasonable epistemic classifications: the type identified must have a reliability that is plausibly correlated with the justificational status of the beliefs in question.

The source of the trivialization is easy to identify. Beebe requires relevant type to be homogeneous with respect to the production of true belief. Suppose that T is a type whose instances include some true belief-producing tokens and some false belief-producing tokens and that its reliability is neither 1 nor 0. Then there is a statistically relevant way to divide T into subtypes: the true belief-producing tokens on one side (T^+) and the false belief-producing tokens on the other (T^-). The subtypes will catapult the reliability to 1 and 0, respectively, so they will be statistically relevant. So T is not homogeneous.

Note that, as pointed out before, if T satisfies the tri-level condition, any subtype of T satisfies it as well. Hence the tri-level condition gives us no grounds to set aside T_+ and T_- as relevant subtypes of T .

We will now consider three amendments of Beebe's proposal that avoid the trivialization. The first tightens the notion of statistical relevance, so that T_+ and T_- are not automatically relevant when T is neither perfectly reliable nor perfectly anti-reliable. The second and third loosen the notion of homogeneity, so that T may count as homogeneous despite the presence of subtypes T_+ and T_- . As we will see, no option is without its difficulties.

5. Tightening statistical relevance

Beebe's approach counts as statistically relevant to the production of true belief any factor that affects the degree of reliability of a type. On reflection, the approach may seem too fine-grained. It would indeed be the way to go if we were interested in the precise degree of reliability pertaining to a given type. But if the issue is whether reliabilism counts a given belief as *justified*, or as *constituting knowledge*, then it suffices to determine whether or not the process was reliable in an all-or-nothing sense, i.e. to determine whether the reliability of the process exceeds a certain threshold value separating processes that are reliable from those that are unreliable. That might not be the only way to think of an absolute property of reliability, but it is certainly the proposal that first comes to mind. Once that matter has been settled, the precise *degree* of reliability is of no concern.

What this suggests is a revised conception of relevance. We may say that a condition C is relevant *with respect to a reliability threshold r* just in case assuming C decreases the

probability of B (true belief) below r or, alternatively, raises the probability of B (true belief) to r or beyond. In more formal terms, C is *r-relevant* to property B in context A just in case: (1) $P(B | A) \geq r$ and $P(B | A \& C) < r$, or (2) $P(B | A) < r$ and $P(B | A \& C) \geq r$. As a special case, where $A \& C$ is a subtype of a belief-forming process type A , C is *r-relevant* to the property of producing true belief just in case either $P(t \text{ produces a true belief} | t \text{ belongs to } A) \geq r$ and $P(t \text{ produces a true belief} | t \text{ belongs to } A \& C) < r$, or $P(t \text{ produces a true belief} | t \text{ belongs to } A) < r$ and $P(t \text{ produces a true belief} | t \text{ belongs to } A \& C) \geq r$.

This is a stronger concept of relevance: a condition that is *r-relevant* is also relevant. The converse does not hold: something may be relevant without being *r-relevant*. We can now define a similarly relativized conception of homogeneity – *r-homogeneity* – with respect to true belief. A subclass S of A is *r-homogeneous* with respect to true belief if and only if there are no *r-relevant* conditions for the occurrence of true belief under circumstances S . Finally, it is tempting to require that the relevant process type for t (with respect to r) be a largest *r-homogeneous* subclass of A (with respect to true belief) within which t falls.

It is easy to see that the revised statistical rule blocks our original trivialization proof. There we showed that a belief-forming process type T would not be homogeneous if its reliability was neither 1 nor 0. For its subtypes T_+ and T_- have reliability 1 and 0, respectively, and they would be statistically relevant in context T because their reliability would then differ from that of T . That step is now blocked, since we cannot assume that T_+ and T_- will be statistically *r-relevant* in context T . So we cannot assume that T is not *r-homogeneous* if its reliability is neither 0 nor 1.

Unfortunately, a similarly trivial conclusion can still be derived. Let t be some token belief-forming process. Suppose T is a type of t selected by the revised constraint: it

satisfies the tri-level condition and it is the broadest r -homogeneous type within which t falls, where r is the threshold for absolute reliability. Suppose first that t produced a true belief. If T is not reliable, then the reliability of T is below threshold r . Trivially, the reliability of $T+$ is above r . So $T+$ is r -relevant in context T and T is not r -homogeneous. It follows that, if t produced a true belief then T is reliable. Similarly, we can show that if t produced a false belief, then T is not reliable.

The source of the problem is also easy to see. If T is not reliable, but has some true belief-producing instances, then it has a subtype that catapults the reliability level above threshold: namely, $T+$. Similarly, if T is reliable, but has some false belief-producing instance, then it has a subtype that lowers the reliability level below the threshold: namely, $T-$.

One may consider another way of coarsening statistical relevance. On Beebe's original proposal, a factor is statistical relevant if it affects a type's reliability in the slightest manner. That seems overly generous. For instance, Beebe assumes the probability of certain types of visual process leading to true belief given that they occur in middle-aged men is *exactly* the same as the probability that they lead to true belief given that they do not occur in middle-aged men (p. 189). If it was not, Beebe would have to count the property of occurring in middle-aged men as part of the relevant process type. But again, the assumption seems excessively strong. A more moderate option would be to say that a factor is relevant if it *significantly* affects the reliability of a process type. More formally, we say that a factor C is *relevant with k significance* to property B in context A if and only if $|P(B | A \& C) - P(A | B)| > k$. A type is k -homogenous with respect to B if it has no subtype that is relevant with k significance. Beebe's notion of relevance corresponds to the limiting case where $k = 0$. On that choice, any slight change in reliability makes for

relevance. By contrast, with $k = .1$, for instance, a factor must raise or lower the reliability by more than $.1$ in order to count as relevant.

The amendment blocks the previous derivation which appealed to the following step: if T is unreliable but contains true beliefs, then T_+ is an r -relevant subtype of T . Hence a true belief-producing token t cannot belong to an r -homogeneous type T which is not reliable. By contrast, we cannot assume that if T is unreliable but contains true beliefs, then T_+ is a significantly relevant subtype of T . For consider a case in which the threshold for absolute reliability is $.9$, the threshold for significance is $.2$, and the reliability of T is $.8$. Since the reliability of T is below $.9$, T is unreliable. However, since the difference in reliability between T and its subtype T_+ is just $.2$, T_+ is not a significantly relevant subtype.

The trivialization can still be derived, however. For this, we assume that a reasonable value for k is below $.5$. The assumption is a mild one: it would be excessive to screen out any factor that affects reliability by less than a half. Suppose that a token true belief belongs to a type T that also contains token processes producing false beliefs. The reliability of T is either lower than or equal to $.5$, or higher than $.5$. In the first case, the difference between the reliability of T and that of T_+ is at least $.5$, so T_+ is a k -significant subtype. In the second case, the difference between the reliability of T and that of T_- is at least $.5$, so T_- is a k -significant subtype. Either way, T is not k -homogeneous. So, if t belongs to a k -homogeneous type, that type must not contain false beliefs. Again, reliability collapses into truth.

6. Loosening homogeneity: no reference to the target outcomes

The trivialization arises because for every type T that includes both true and false beliefs, $T+$ and $T-$ undermine the homogeneity of that type. As the result, the homogeneity condition is only satisfied by perfectly reliable and perfectly anti-reliable types. The trivialization follows. To avoid it, we need to rule out subtypes such as $T+$ and $T-$ as admissible restrictions of T .

A first way to do so follows closely Salmon's (1971) solution to the *reference class problem*. While Beebe's statistical requirement is based on Salmon's idea, Beebe unfortunately left out a crucial *provisio* in the definition of homogeneity. The *provisio* avoids straightforward trivialization. But it faces related difficulties. Let us explain.

On a frequentist notion of probability, we cannot assign non-trivial probabilities to single events unless we assign these events to some class. But a single event belongs to an infinity of classes. Hence we need some way of selecting a particular class as the relevant one. That is the reference class problem. It is, of course, strongly reminiscent of the generality problem. Salmon's solution states that the probability of a single event e to produce an outcome of type B is given by the class containing e which is the *broadest homogeneous one with respect to the outcome B* . A class A is homogeneous with respect to B iff there are no *admissible* subtypes of A that are statistically relevant to B . Thus, for any *admissible* subtype $A\&C$ of A , it should be the case that $P(B | A\&C) = P(B | A)$.

Without a restriction to admissible subtypes, Salmon's solution would be trivial. For suppose $0 < P(B | A) < 1$. Then there is a subtype of A , namely $A\&B$, such that $P(B | A\&B) = 1 \neq P(B | A)$. So A is not homogeneous with respect to B . All single event probabilities would collapse into 0 or 1. The collapse of Beebe's proposal is a special case of this more general problem. But Salmon does not allow any subtype to count against homogeneity. Here is what he has to say:

A place selection [i.e., a subtype] effects a partition of a reference class into two subclasses, elements of the place selection and elements not included in the place selection. In the reference class of draws from our urn [containing red, blue and white balls], every third draw starting with the second, every k th draw where k is prime, every draw following a red result, every draw made with the left hand, and every draw made while the sky is cloudy all would be place selections. *"Every draw of a red ball" and "every draw of a ball whose color is at the opposite end of the spectrum from violet" do not define place selections, for membership in these classes cannot be determined without reference to the attribute in question.* (Salmon 1971, p. 43, our emphasis)

Hence, admissible subtypes are those that can be “determined without reference to” the outcome we are considering. For instance, the statistical relevance of the subtype $A\&B$ cannot count against the homogeneity of A with respect to outcome B , because the former cannot be determined without reference to the outcome B . The *provisio* is essential to Salmon’s definition, but left out by Beebe (cf. the characterizations of homogeneity in Beebe 2004, pp. 181, 188-189).

The revised conception of homogeneity bars trivialization. Suppose T is a type for a true-belief producing token t that is the broadest homogeneous type of t with respect to the outcome of producing true beliefs. If T is less than perfectly reliable, then T_+ is a statistically relevant subtype of T . But T_+ is not an *admissible* statistically relevant subtype of T , for T_+ is characterized by reference to the target outcome, namely, producing true beliefs. And similarly for T_- .

But this is not the end of the story. For while we have ruled out subtypes that are “determined with reference to” the outcome of true belief, there will be many subtypes that involve a factor that is (a) irrelevant to the justification of the belief, (b) statistically

relevant, (c) not in any reasonable sense determined with reference to the outcome. The following example illustrates the problem.

Let us return to Smith's story and stipulate that Smith's procedure for visually identifying maple trees is a very unreliable one. For instance, let us assume that Smith classifies any tree whose leaves are biggish as a maple tree. (More realistic examples can be given.) Intuitively, the relevant type for evaluating Smith's belief is that procedure, possibly supplemented with an information-processing problem (to identify the species of a visually perceived nearby tree, perhaps), details of the algorithm (perhaps including some functional description of Smith's visual processing) and of Smith's "cognitive architecture". And intuitively again, the type should come out as unreliable, and Smith's belief as unjustified.

We strongly doubt that the revised proposal can deliver the desired verdict. Let T be our desired type. Consider the property of "occurring in the vicinity of one or more maple trees" (M). Plausibly, the probability of acquiring a true belief by an instance of T in the vicinity of maple trees is much higher than the probability of acquiring it by an instance of T . That is so because Smith's procedure will qualify pretty much any maple tree as a maple tree, but most other trees as maple trees as well. So $T\&M$ is a statistically relevant subtype of T with respect to the outcome of producing a true belief. Worse, $T\&M$ might well be a reliable subtype. Yet given our stipulations, $T\&M$ is not relevant to the justification of Smith's belief and Smith's belief is not justified.

To deliver the right verdict, defendants of the revised proposal would need to rule out subtypes like $T\&M$ as inadmissible. Could they do so? The answer is not straightforward, because the notion of a type "being determined by reference to the outcome" is not clear. The defenders would need to say more about it. Still, *prima facie* at least, "occurring in the vicinity of maple trees" does not obviously make reference to the fact that Smith's

tree-identification process produces a true belief. Granted, in the case at hand, that feature significantly increases the *probability* that Smith's process produces a true belief. But we certainly do not want to rule out any feature that increases Smith's chances of acquiring a true belief. In that respect, "occurring in the vicinity of maple trees" does not differ from "occurring in a high level of ambient light", which Beebe is happy to admit (p. 189). Moreover, even if we conceded that "occurring in the vicinity of maple trees" does make reference to the process's outcome, there are many other features that could less easily be claimed to do so, but that are equally irrelevant to the justification of Smith's belief and yet statistically relevant to his forming a true belief. Candidates are, for instance, "occurring in an area where there were maple trees a few seconds ago", "occurring while perceiving an object that many people take to be a maple tree", or "occurring in the vicinity of bits of DNA of type *X*", where *X* is a chemical description of DNA that is maple-tree specific, and so on.

Beebe and defenders of the revised conception could take issue with the details of the example. They may argue that we haven't specified Smith's "information-processing problem", "algorithm" or "cognitive architecture" properly. On an adequate specification the epistemically irrelevant factors we put forward will perhaps turn out to be statistically irrelevant. In reply, we first note that the rejoinder underlines the uncertain application of the notions of "information-processing problem", "algorithm" or "cognitive architecture" themselves, thus raising the suspicion that Beebe's solution is partly question-begging in ignoring a generality problem that arises in the application of these terms themselves. (We return to the issue in section 8.) Second, we think the example illustrates a structural problem. Rejecting the details of the example is not enough to convince us that the problem does not arise.

The structural problem is this. The core idea of Salmon's solution to the reference class problem for single event probability is to take into account *all the available relevant factors*. By "available", we mean factors that are in some sense fixed "before" or "independently of" the outcome. By "relevant", we mean that they affect the probability of the outcome. As a result, the relevant class is *as specific as needed* to take into account all such factors, but *not more*. This proposal makes sense when ascribing objective probabilities to single case events. But when applied to the production of true belief, it enjoins us to take into account all information that is relevant to whether one will produce a true belief, as long as it is available in some sense "before" the belief is produced. That is why the presence of maple trees nearby, their presence nearby a minute ago, the presence of maple-tree DNA around, and so on, are taken into account. The procedure results in something that can properly be called *the probability that Smith forms a true belief in the present circumstances*, thus delivering a single-event probability. But that probability is typically epistemologically irrelevant: if circumstances are exceptionally favorable, the probability will be high even though Smith's belief is (intuitively) unreliably formed.

In a nutshell, by taking over Salmon's idea, Beebe in effect identifies the reliability of a token belief with the single-case probability that it is true. But, once more, the epistemological status of beliefs does not correlate with the single-case probability that they are true. The problem arises however we spell out the notions of a type being "determined with reference to the outcome", "information-processing problem", "algorithm" or "cognitive architecture". Hence Beebe's proposal as well as the revised version we have considered here are bound to fail.

7. Loosening homogeneity: natural kinds only

Another way to avoid the trivialization is to restrict admissible types to natural kinds, where a natural kind is thought of as a kind that appears substantially in some natural law. For a natural type T of belief-forming process, there is little reason to think that $T+$, its conjunction with the property of producing a true belief, is itself a natural type. So the revision will equally prevent the likes of $T+$ and $T-$ from counting against homogeneity.

In fact, the statistical approach seems well-suited for solving the specificity problem facing Alston's theory. As we recall, Alston advocates choosing as the type of a given token process the most specific psychologically realistic class. We also saw that this will, in Conee and Feldman's view, result in too specific a categorization of token processes. Now Beebe's restriction to types that satisfy the tri-level condition is roughly equivalent to Alston's restriction to psychologically realistic process types. But according to the statistical approach, once again, the relevant type among those is not the most specific one, but the broadest statistically homogeneous one. That the type is homogeneous means that no admissible subtype is statistically relevant to the production of true belief. The homogeneity condition is by itself sufficient to prevent the statistical proposal from collapsing into Alston's idea of maximum specificity. While the maximally specific type will also be statistically homogeneous (in a trivial sense), the converse is not true in general: statistically homogeneous types will not in general be maximally specific. Add to this the further suggestion that we are to choose the *broadest* statistically homogeneous type that satisfies the tri-level condition and it becomes clear that Beebe's statistical rule, if taken to operate on natural kinds, avoids maximal specificity and what Feldman calls the single case problem.

The natural kind restriction may also provide some leverage against the specificity problem raised in the previous section. Take the example of unreliable Smith's again, and let T be our desired type. One may hope that subtypes such as "instance of T occurring in

the vicinity of maple trees” do not correspond to natural kinds. If so, their presence does not count against *T* being homogeneous and the previous objection would be blocked.

Even though the statistical approach looks quite promising when applied to natural kinds, a problem remains that has its roots in scientific diversity. As Conee and Feldman’s observe, different sciences adopt different classificatory schemes and different partitions of the world into natural kinds (Conee and Feldman 1998, pp. 10-11). The concept of a natural kind is plausibly construed, therefore, as being relative to a given science. Yet there is no guarantee that the categories in one science are extensionally equivalent to the categories of another. Instead, things that count as being of the same type in one science may be classified as belonging to different types in another. In Alston’s view, we should focus on psychological kinds, but – again following Conee and Feldman – it is not at all clear why psychology, as opposed to neurology or chemistry, should be granted a privileged status.¹⁰

There are essentially three responses to this worry. First, we could group together all the different scientific classificatory schemes into one superclassification. This should be done, most naturally, by adopting a conceptual framework that admits all possible scientific distinctions to be made. (If in doubt whether *all* sciences are relevant, the reader may select a subset of sciences that she thinks are.) If one science allows a distinction between *As* and $\sim A$ s and another one between *Bs* and $\sim B$ s, the superclassification would contain the types $A\&B$, $A\&\sim B$, $\sim A\&B$ and $\sim A\&\sim B$. For example, *A* could be the class of all psychological processes of a certain kind and *B* the class of all neurological processes of a certain kind. $A\&B$ would then correspond to a type of process that has both the psychological *A*-characteristics and the neurological *B*-characteristics. The problem is that we have no assurance that the categories in the superclassification are themselves

¹⁰ The conceptual diversity of science is stressed in Dupré (1993).

natural kinds, for there is no guarantee that a given class in the superclassification itself figures substantially in any important scientific generalization.

The other response to this worry would be to apply the statistical approach to each scientific classification separately and then somehow combine the resulting broadest homogeneous types into one single type. But there is, for similar reasons, no guarantee that the resulting “interdisciplinary” type will itself be a natural kind. A third reaction might involve arguing for a reductive thesis according to which all of science can be reduced to one fundamental science, most plausibly some part of physics, and that the categories adopted by that science are the *real* natural kinds. However, it is difficult to see how a purely physical description of a belief forming process could have any bearing on the question of whether processes of that type generally lead to true beliefs. Unsurprisingly, no well-known approach to the generality problem advocates a purely physicalistic account of process types.

Even if can live with the fact that the relevant classification of a given process token, although constructed from natural kinds, may itself fall short of being a natural kind, there is still the worry of uniqueness, for we may ask what reason there is to think that a *unique* broadest homogeneous type will result from applying the statistical rule to natural kinds. Why could there not be several maximally broad homogeneous natural kind based types? Suppose that *C* is statistically homogenous (with respect to true belief) but not broadest. Nothing seems to prevent that “*A* or *C*” and “*B* or *C*” are both statistically homogeneous, but that no superset of either set is, in which case both are “broadest”.

8. Further technical issues

Before we conclude, we would like to register some further technical issues with Beebe's proposal, including the uniqueness worry raised in section 3.

First, Beebe assumes that each token belief-forming process belongs to at least one type that satisfies the tri-level condition, and that there is a broadest such type (cf. "Let A be the broadest process type that satisfies the tri-level condition for some process token t ", Beebe 2004, p.180, pp. 187-188). If we grant Marr's tri-level hypothesis, the first assumption is tantamount to saying that each token belief-forming process is a cognitive process. That is fairly unproblematic: if there are counterexamples, the corresponding beliefs can presumably be classified as unjustified.

The assumption that there is a *broadest* such type raises much more serious issues, and Beebe does not argue for it. The assumption is equivalent to the following: if a token process t belongs to two types T and T' satisfying the tri-level condition, then there must be some type T'' (possibly the same as T or T') that satisfies the tri-level condition and that includes all instances of T and T' . We can see two ways to legitimize the assumption. The first is to claim that (a) for any information-processing problems i and i' , there is an information-processing problem i'' consisting in solving i or i' , (b) for any algorithms m and m' , there is an algorithm m'' consisting in following m or m' , and (c) for any cognitive architectures s and s' , there is a cognitive architecture s'' consisting in having the fundamental capacities in s or s' . The claim conflicts with Beebe's use of the tri-level condition to discard such types as "visual process". For it entails that the collection of information-processing problems, algorithms and architectures involved in visual processes constitute a unique (massively disjunctive) information-processing problem, a unique (massively disjunctive) algorithm and unique (massively disjunctive) cognitive architecture, so that the type satisfies the tri-level condition after all.

Another way to defend the assumption is to claim that for a given token process t , there is a *unique* information-processing problem that it solves, a *unique* algorithm that it implements, and a *unique* cognitive architecture on which the algorithm is executed. Then the broadest type is well-defined: it consists in *solving the information-processing problem solved by t following the algorithm followed by t on the cognitive architecture of t* . Bebe's use of definite descriptions in statement of his tri-level condition suggests that he endorses this claim (cf. quote in section 2 above). However, the claim is likely to be seen as question-begging in the context of the generality problem. Part of the generality problem is that the token process that leads to Smith's belief can be seen as answering several information-processing problems: identifying the species of a tree that one sees, identifying the species of that particular tree, telling whether there are maple trees nearby, and so on. Similar worries arise for the idea that there is a unique algorithm and a unique cognitive architecture that this process and Smith instantiate. By assuming that a token process will select a particular type of tri-level description, Beebe is in effect presupposing a partial solution to the problem. We thus fail to see how Beebe can unproblematically assume that each belief-forming token is associated with a unique broadest type satisfying the tri-level condition.

Second, Beebe assumes that given a type A of t satisfying the tri-level condition, there is a unique broadest homogeneous type to which t belongs. Thus he writes:

If A [the broadest type of t satisfying the tri-level condition] is partitioned into a set of maximally and objectively homogeneous subclasses $A\&C_1, \dots, A\&C_n$, the value of $P(t' \text{ produces a true belief} \mid t' \text{ belongs to } A\&C_j)$ will be the same for every t subsumed by each $A\&C_j$. Each cell in such a partition will be the relevant process type for every process token that falls under it. (Beebe, p.190, notation adapted)

The assumption does not hold in general. Suppose that A can be divided according to two features, B and C , and only according to them. Let $P(t \text{ produces a true belief} \mid t \text{ belongs to } A) = 3/4$, and let the conditional probabilities of true belief given subtypes of A be as follows:

	$A \& C$	$A \& \sim C$
$A \& B$	1	1
$A \& \sim B$	1	0

For instance, you may suppose that each cell has only one instance, and that probabilities are assigned according to the number of instances: in the first three cases a true belief, in the case of $A \& \sim C \& \sim B$ a false one. What is the broadest homogeneous type of t ? A itself is not homogeneous: its reliability is $3/4$, but it has statistically relevant partitions, such as $A \& C$ and $A \& \sim C$. $A \& B$ and $A \& C$ are homogeneous: $A \& B$ can be divided into $A \& B \& C$ and $A \& B \& \sim C$. But this partition is not statistically relevant. The same goes for $A \& C$. $A \& B \& C$ is homogeneous because it cannot be further subdivided, but it is not broadest, since it is a subset of $A \& B$ which is homogeneous. Finally, $A \& B$ and $A \& C$ are both homogeneous and broadest, but none is included in the other. So there is no unique broadest homogeneous subtype.

In short, Beebe assumes that “repeated” division of a type into statistically relevant subtypes must converge on a unique subtype. The example above shows that that is not so. If we start to divide A into $A \& B$ and $A \& \sim B$, we will stop at $A \& B$. If we start instead by dividing it into $A \& C$ and $A \& \sim C$, we will stop at $A \& C$. If we start by dividing into $A \& B \& C$, $A \& B \& \sim C$, $A \& \sim B \& C$ and $A \& \sim B \& \sim C$, we will reach a homogeneous type that is

not broadest, but we will not know which way to enlarge it. Thus, Beebe's solution succumbs to the uniqueness worry we raised earlier.

9. Conclusion

The purpose of this paper was to consider, in some depth, the promising statistical approach to the generality problem due to James R. Beebe. The approach has the virtues of addressing the generality problem directly in a manner that is exact, principled and true to the spirit of reliabilism. However, our results were disappointing. As we just showed, the statistical approach may fail to select a unique relevant type. Even if we put this more technical problem aside, the strategy leads to a disturbing result: if it succeeds at all in picking out a type for a given token process, then reliability collapses into truth (and anti-reliability into falsehood). There was seen to be no quick fix to this problem. The threat of trivialization reappears if we relativize the notion of statistical relevance to a reliability threshold or to a level of significance. Other problems arise if, instead, we restrict subtypes to natural kinds or to ones that can be picked out without reference to the truth of the belief in question. The conclusion is inevitable: reliabilists have to look elsewhere for a convincing solution to the generality problem.

References

- Adler, J. (2005), "Reliabilist Justification (or Knowledge) as a Good Truth-Ratio", *Pacific Philosophical Quarterly* 86: 445-458.
- Adler, J., and Levin, M. (2002), "Is the Generality Problem too General?", *Philosophy and Phenomenological Research* 65 (1): 87-97.

- Alston, W. (1995), "How to Think about Reliability", *Philosophical Topics*, Spring: 1-29.
- Baumann, P. (2011), "A Puzzle about Responsibility: A Problem and its Contextualist Solution", *Erkenntnis* 74(2): 207-224.
- Beebe, J. R. (2004), "The Generality Problem, Statistical Relevance and the Tri-Level Hypothesis", *Nous* 38: 1, 177-195.
- Bishop, M. A. (2010), "Why the Generality Problem is Everybody's Problem", *Philosophical Studies* 151: 285-298.
- Brandom, R. B. (1998), "Insights and Blindspots of Reliabilism", *The Monist* 81 (3): 371-392.
- Christensen, D. (2007), "Three Questions about Leplin's Reliabilism", *Philosophical Studies* 134 (1): 43-50.
- Comesaña, J. (2006), "A Well-founded Solution to the Generality Problem", *Philosophical Studies* 129: 27-47.
- Conee, E. and Feldman, R. (1998), "The Generality Problem for Reliabilism", *Philosophical Studies* 89: 1-29.
- Dupré, J. (1993), *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*, Harvard University Press.
- Feldman, R. (1985), "Reliability and Justification", *The Monist* 68: 159-174.
- Feldman, R., and Conee, E. (2002), "Typing Problems", *Philosophy and Phenomenological Research* 65 (1): 98-105.
- Goldman, A. I. (1979), "What is Justified Belief?", in Pappas, G. (ed.), *Justification and Knowledge*, Dordrecht: D. Reidel.

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

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

Goldman, A. I. (2008), "Reliabilism", *The Stanford Encyclopedia of Philosophy* (Summer 2008 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2003/entries/davidson/>.

Goldman, A. I., and Olsson, E. J. (2009), "Reliabilism and the Value of Knowledge", in Haddock, A., Millar, A. and, Pritchard, D. H. (eds.), *Epistemic Value*, Oxford University Press.

Heller, M. (1995), "The Simple Solution to the Problem of Generality", *Noûs* 29 (4): 501-515.

Hetherington, S. C. (1996), *Knowledge Puzzles: An Introduction to Epistemology*, Boulder, CO: Westview.

Kappel, K. (2006), "A Diagnosis and Resolution to the Generality Problem", *Philosophical Studies* 127: 525-560.

Lemos, N. (2007), *An Introduction to the Theory of Knowledge*, New York: Cambridge University Press.

Leplin, J. (2007), "In Defense of Reliabilism", *Philosophical Studies* 134 (1), 2007: 31-42.

Lycan, W. G. (1988), *Judgment and Justification*, New York: Cambridge University Press.

Marr, D. (1982), *Vision*, San Francisco: W. H. Freeman.

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., and Cruz, J. (1999), *Contemporary Theories of Knowledge*, 2nd edition, Lanham, MD: Rowman and Littlefield.

Salmon, W. C. (1971) "Statistical Explanation," in W. C. Salmon, R. Jeffrey, & J. Greeno, *Statistical Explanation and Statistical Relevance*, Pittsburgh: University of Pittsburgh Press.

Ramsey, F. P. (1931), *The Foundations of Mathematics and other Logical Essays*, Braithwaite, R. B. (Ed.), London: Routledge and Kegan Paul.

Wunderlich, M. (2003), "Vector Reliability: A New Approach to Epistemic Justification", *Synthese* 136: 237-262.