

A New Statistical Solution to the Generality Problem

Samuel Kampa

[Penultimate draft. Forthcoming in Episteme. Please cite published draft.]

1. Introduction

The Generality Problem is the problem of identifying, for any given belief, which belief-forming process type is relevant for justification-determining purposes. First noted by Alvin Goldman (1979, p. 39), the problem was initially presented as a serious challenge for process reliabilism. In recent years, it has been suggested that the Generality Problem is a problem for every major theory of justification, not just reliabilism.¹ If the Generality Problem is everyone's problem, finding a solution to the Generality Problem is in everyone's interests. It should come as a surprise, then, that the Generality Problem has received somewhat scant attention. Few epistemologists write articles on the subject, and even fewer hazard solutions.² And while the solutions currently on offer may ameliorate the problem, they aren't ideal. An ideal solution is principled, rigorous, sparing in the conceptual resources it introduces, plausible in its account of cognitive mechanisms, and (of course) possibly true.³ No extant solution satisfies each of these desiderata. Despite the sobering track record, however, I maintain that there is in fact an ideal solution, and I intend to provide it.

My solution builds upon the closest thing we have to an ideal solution, namely, James R. Beebe's (2004) "Statistical Solution". Beebe's solution satisfies nearly all of the foregoing desiderata. There's just one desideratum it lacks: being possibly true. In their sustained treatment of the Statistical Solution, Julian Dutant and Erik Olsson (2013) successfully refute it. They go a whit too

¹ See Comesaña (2006) and Bishop (2010). See also Conee's (2013) response to Bishop.

² Perhaps no treatment of the Generality Problem is as impressive as that of Conee and Feldman (1998). Solutions have been proposed by Goldman (1986, pp. 49-51), Schmitt (1992), Alston (1995), Wunderlich (2003), Beebe (2004), Becker (2008), and Olsson (2016). But as Dutant and Olsson (2013) note: "although much work has gone into solving the Generality Problem, none of the proposals that have emerged so far has gained widespread acceptance" (p. 1351). While a full treatment of the aforementioned solutions is beyond the scope of this paper, certain of their shortcomings may become apparent along the way.

³ Feldman (1985) and Conee and Feldman (1998) demand as much from a solution to the Generality Problem. Dutant and Olsson (2013) concede that Beebe's solution is in many ways ideal, but they nonetheless abandon the statistical approach for reasons we consider in section 4.

far, however, in maintaining that all Statistical Solutions are dubious. On the contrary, there is a successful Statistical Solution that removes the defects and retains the virtues of Beebe's solution. This, at any rate, is what I aim to show.

I begin by providing a very broad introduction to the Generality Problem. I then rehearse Beebe's solution and Dutant and Olsson's most important objection, and I add an objection of my own. Finally, I advance a solution that suffers no such defects: what I call the *New Statistical Solution*.

2. The Generality Problem

The Generality Problem was initially presented as a problem for reliabilist theories of justification.⁴ According to reliabilism, a belief is justified if and only if it is produced by a reliable belief-forming process. More precisely, reliabilism says that a belief is justified if and only if it is produced by a belief-forming process token whose unique relevant belief forming process type is reliable. Take the process token that produces the belief that I am currently writing on the Generality Problem. According to reliabilism, this process token belongs to one and only one relevant process type.⁵ Whether this process token produces a justified belief depends on whether the tokened unique relevant type is reliable.

We can present the Generality Problem in the form of a question (call it the *Generality Question*): *how do we determine which process type is relevant for a given process token?* To appreciate this question's force, let us again consider my belief that I am currently writing on the Generality Problem. The process token that produced this belief falls under a broader process type: *inference*.⁶ But it also falls under more specific types, including *inferring under normal conditions*, *inferring that one is*

⁴ Alternatively, some frame the Generality Problem as problem for reliabilist theories of *knowledge*. For our purposes, nothing of interest hangs on which formulation we choose.

⁵ Most reliabilists hold that there must be only one relevant type for each token process. For an alternative reliabilist view, see Wunderlich (2003).

⁶ One might think that the token under consideration in fact falls under a different type (perhaps *direct acquaintance*). We needn't settle the question here. For our purposes, virtually any broad process type will do.

writing a paper, inferring on a Wednesday, and so on. The question is: which (if any) of these process types is relevant for determining whether and to what extent my belief is justified?

However we answer the question, we must avoid two extremes: picking too specific a process type, and picking too general a process type. These extremes correspond to the *Single Case Problem* and the *No-Distinction Problem*, respectively:⁷

Single Case Problem: For each process token t , the process type T picked out as relevant has only t for a token. As a result: *a)* T is unrepeatable, and *b)* T is either perfectly reliable or perfectly unreliable.

No-Distinction Problem: For each process token t_1, t_2, \dots, t_n , the process type T picked out as relevant for t_1, t_2, \dots, t_n fails to distinguish adequately between differing degrees of justification in the beliefs produced by t_1, t_2, \dots, t_n .

Consider an instance of the Single Case Problem. Say that one picks out *inferring that one is writing this essay on January 1 at noon* (T_1) as the relevant process type for process token t_1 (where “this” is treated as a rigid designator). T_1 has only one process token (t_1), and t_1 produces a true belief.⁸ The problem that, since T_1 ’s only token produces a true belief, T_1 is perfectly reliable, and hence, the belief produced by t_1 is perfectly justified. But surely my belief that I’m currently writing this essay is not, in fact, perfectly justified. Moreover, for reliabilism to be at all interesting, process types must be *shareable*, which is to say that they must possibly have multiple tokens. But T_1 isn’t sharable; it is a “one and done” type. This we cannot allow. A good answer to the Generality Question must rule out single cases.

⁷ Feldman (1985)

⁸ Technically, T_1 , as it is described here, is repeatable. A more fully described version of T_1 would include the year, millisecond, etc., at which the inference took place. For ease of reading, however, I will stick with the abridged description of T_1 .

Consider an example of the No-Distinction Problem. Say that one picks out *inferring* (T_2) as the relevant type for process token t_1 . T_2 has among its tokens: *the inference that one is writing a paper on January 1 at noon* (t_1), *an instance of hasty generalization* (t_2), *an instance of inferring the antecedent from the consequent* (t_3), and so on. It is obvious that the belief t_1 produces is better justified than the beliefs t_2 and t_3 produce. But if T_2 is the unique relevant type to which t_1 , t_2 , and t_3 belong; and if, for each belief b and token t , if t produces b , b 's degree of justification is a direct function of the degree of reliability of the unique relevant type to which t belongs; then t_1 , t_2 , and t_3 yield equally justified beliefs. This is a bad result—one that emerges from selecting a process type that is too broad to countenance genuine differences in justification between beliefs. Thus, a good answer to the Generality Question must also rule out no-distinction cases.

Conee and Feldman (1998) propose three additional criteria for a good answer. First, a good answer must be principled. No arbitrary rules or *ad hoc* stipulations will do. Second, it must “make defensible epistemic classifications” (p. 4). Not just any solution to the No-Distinction Problem will suffice; rather, a good answer countenances genuine differences in justification between beliefs. Third, it “must remain true to the spirit of the reliabilist approach” (ibid.). Rules that smack of internalism thus deserve the boot.

The challenge, then, is to devise a solution that satisfies the above criteria and dodges both the Single Case Problem and the No-Distinction Problem. It is a challenge that reliabilism is highly unlikely to meet, according to Conee and Feldman. They pronounce, with histrionic intensity, that “...the prospects for a solution to the Generality Problem for process reliabilism are worse than bleak...In the absence of a brand new idea about relevant types, the problem looks insoluble. Consequently, process reliability theories of justification and knowledge look hopeless” (pp. 5, 24).

Is the outlook for reliabilism really so bleak? Might there still be a good answer to the Generality Question? Purported solutions have come and gone, but as Dutant and Olsson rightly

note: “none of the proposals that have emerged so far has gained widespread acceptance” (2013, p. 1351). This is partly due to the fact that most extant solutions either come packaged with additional substantive theoretical commitments⁹ or solve only one version of the Generality Problem¹⁰ or break with orthodox reliabilism¹¹ or introduce new problems.¹² But Beebe’s solution is different. It involves minimal theoretical commitments, purports to solve all versions of the Generality Problem, requires no controversial modifications to reliabilism, and introduces no serious new problems. Moreover, it is principled, rigorous, and strikingly elegant. It appears, then, that if any approach to the Generality Problem holds promise, it is some version of the Statistical Solution. In what follows, I consider Beebe’s Statistical Solution, discuss its defects, and propose a New Statistical Solution that suffers no such problems.

3. Beebe’s Statistical Solution

I cannot, in the space provided, do justice to the nuances of Beebe’s account. While the exposition that follows is unavoidably terse, I trust that it at least conveys the most important details of his account.

Beebe’s solution is presented in two parts. First, he provides a partial definition of process types. Second, he proposes a rule for picking out relevant process types. The first part yields the *Tri-Level Condition*; the second yields the *Statistical Answer*. I begin by presenting and expounding the first condition:

⁹ e.g., Becker (2008), Comesaña (2006)

¹⁰ e.g., Olsson (2016)

¹¹ e.g., Wunderlich (2003)

¹² e.g., Comesaña (2006), Schmitt (1992). See also Conee and Feldman (1998, pp. 18-20) on Schmitt.

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 ” (p. 180).

Here, Beebe is taking a cue from neuroscientist David Marr (1982), who formulated an early version of the *Tri-Level Hypothesis*. The main idea is that a complete description of a cognitive system requires explanations at the computational level, the algorithmic level, and the implementation level. Computation-level explanations answer the question: “What information-processing problem is the system in question solving?”; algorithm-level explanations answer the question: “What method is the system using to solve this information processing problem?”; and implementation-level explanations answer 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). While Marr’s more detailed account of visual perception is somewhat dated, the Tri-Level Hypothesis remains influential as a high-level conceptual paradigm in cognitive science.¹³

Beebe’s Tri-Level Condition partially individuates process types by their computational and algorithmic properties.¹⁴ By divvying up types in this way, Beebe aims to “significantly [reduce] the field of potentially relevant cognitive types” (p. 180). The Tri-Level Condition rules out most of the

¹³ A full defense of the Tri-Level Hypothesis is beyond the scope of this paper. For more on the Tri-Level Hypothesis, see Beebe (2004, pp. 181-3), Dawson (1998), and Marr (1982). See also Prinz (2007) for an endorsement of a broadly Marrian view, and Prinz (2012, pp. 50-4) for proposed revisions to Marr’s account of visual perception. I am grateful to Julia Haas and Eric LaRock for filling me in on the status of the Marrian paradigm in contemporary cognitive science and philosophy of perception.

¹⁴ Note that implementation-level considerations aren’t built in to the Tri-Level Condition. Since Beebe is dealing with process types, and since one and the same process type can be tokened multiple times by physically distinct cognitive systems, no implementation-level account is necessary. See Beebe (2004, pp. 184-5).

(pseudo-)types that Conee and Feldman (1998) use as examples. As Beebe notes, the examples they give are either *a)* too vague with respect to computational and algorithmic properties, or *b)* too physically descriptive, and thus unduly pitched at the implementation level. So Conee and Feldman fail to provide examples of genuine types—or, at the very least, their examples underdetermine whether the types in question are genuine. The *really* genuine types will have computational and algorithmic properties built into their descriptions. More precisely:

Tri-Level Definition: Process type T is partially defined as the conjunction of I_s , M_p , and S_u , where I_s is the property of solving information problem i_s , M_p is the property of using procedure m_p to solve an information-processing problem, and S_u is the property of executing a problem-solving procedure on a system of type s_u .¹⁵

The Tri-Level Condition and the Tri-Level Definition combine to give the following partial answer to the Generality Problem:

Tri-Level Answer: For any process token t , T is the relevant process type only if *a)* t tokens T , and *b)* for some I_s , M_p , and S_u , T is partially defined as $I_s \bullet M_p \bullet S_u$.¹⁶

The Tri-Level Answer is a good first step. But it's just that: a *first* step. To see why, consider type T and distinct properties F and G , and say that: *i)* T instantiates I_s , M_p , and S_u , and *ii)* t tokens T , $T \bullet F$, and $T \bullet G$. Since by stipulation, T is partially defined as $I_s \bullet M_p \bullet S_u$ and t is a token of T , T satisfies conditions *(a)* and *(b)* of the Tri-Level Answer. The problem is that T isn't the *only* type of t that satisfies *(a)* and *(b)*. Take type $T \bullet F$. Since T is partially defined as $I_s \bullet M_p \bullet S_u$, so is $T \bullet F$; thus, *(a)* is satisfied. And since t tokens $T \bullet F$ by stipulation, *(b)* is satisfied. The same goes for $T \bullet G$. So which

¹⁵For a more detailed statement of the Tri-Level Definition, see Beebe (2004, pp. 183-4). Note: "Tri-Level Definition" is my term, not Beebe's.

¹⁶For more on this, see Beebe (2004, pp. 184, 193).

type is relevant: T , $T \bullet F$, or $T \bullet G$? The Tri-Level Answer doesn't say.¹⁷ In order to determine the relevant type, we need a further condition:

Statistical Answer: For any process token t , T is the relevant process type only if, *c*) for any class \mathcal{A} of which T is a subclass and in which t falls, “ T is the *broadest objectively homogenous subclass* of \mathcal{A} ” (Beebe, 2004, p. 188).

Treating *(a)*, *(b)* and *(c)* as jointly sufficient conditions yields the *Statistical Solution to the Generality Problem*.

Before considering how the Statistical Solution improves upon the Tri-Level Answer, we need to do a bit of unpacking. The following definitions serve as a sufficient albeit dense explication of the Statistical Answer:¹⁸

- i) A subclass T is objectively homogenous if and only if “there are no statistically relevant partitions of T that can be effected” (Beebe, 2004, p. 188).
- ii) Partition K of T is statistically relevant if and only if $P(t \text{ produces a true belief} \mid t \text{ belongs to } T) \neq P(t \text{ produces a true belief} \mid t \text{ belongs to } K)$. Informally, a partition is statistically relevant if it's difference-making with respect to probability.
- iii) T is the *broadest* objectively homogenous subclass of \mathcal{A} if and only if *a*) T is a subclass of \mathcal{A} , *b*) T is objectively homogenous, and *c*) T is not a proper subclass of another objectively homogenous subclass of \mathcal{A} . Informally, T is *maximally* objectively homogenous if it's not part of a larger objectively homogenous subclass.

¹⁷ Beebe (2004, pp. 187-8)

¹⁸ For *(i)*-*(iii)*, assume that t is a token of both K and T .

We are now in a position to appreciate the upshot of the Statistical Answer. Let us return to our previous example, where T , $T \bullet F$, and $T \bullet G$ satisfy the Tri-Level Condition. Which type is relevant for t ? According to the Statistical Answer, it's the broadest objectively homogenous process type containing t , i.e., the largest type-class containing t that doesn't have any statistically relevant partitions. Say that $P(t \text{ produces a true belief} \mid t \text{ belongs to } T) = P(t \text{ produces a true belief} \mid t \text{ belongs to } T \bullet F) = P(t \text{ produces a true belief} \mid t \text{ belongs to } T \bullet G)$. Say further that $T \bullet F$ and $T \bullet G$ are T 's only partitions and that, for every class \mathcal{A} of which T is a proper subclass, $P(t \text{ produces a true belief} \mid t \text{ belongs to } T) \neq P(t \text{ produces a true belief} \mid t \text{ belongs to } \mathcal{A})$. Then T is the relevant type of t ; it's the largest objectively homogenous subclass to which t belongs. Alternatively, say that $P(t \text{ produces a true belief} \mid t \text{ belongs to } T) = P(t \text{ produces a true belief} \mid t \text{ belongs to } T \bullet F) = 0.8$, but that $P(t \text{ produces a true belief} \mid t \text{ belongs to } T \bullet G) = 0.5$. Say further that $T \bullet G$ has no statistically relevant partitions. Then $T \bullet G$ is the relevant type of t ; it's the largest objectively homogenous subclass to which t belongs.

A less abstract example may be of use. Consider the following three process types: inferring (T), inferring on the basis of sense perception under favorable conditions ($T \bullet S$), and inferring on the basis of sense perception under favorable conditions on a Wednesday ($T \bullet S \bullet W$). Now consider a token inference t , and (for simplicity's sake) assume that t tokens all and only T , $T \bullet S$, and $T \bullet S \bullet W$. For the Statistical Solution to be worth its weight, it needs to pick out $T \bullet S$ as the relevant type; T is too broad (and thus succumbs to the No-Distinction Problem), and $T \bullet S \bullet W$ has a cognitively irrelevant factor in its description. And indeed, given certain plausible probability assignments, the Statistical Solution delivers the correct verdict. It is obvious that $T \bullet S$ is more reliable than T . So $P(t \text{ produces a true belief} \mid T \bullet S) > P(t \text{ produces a true belief} \mid T)$, and hence $T \bullet S$ is a statistically relevant partition of T . It is also the case that $T \bullet S \bullet W$ and $T \bullet S$ are equally reliable. So $P(t \text{ produces a true belief} \mid T \bullet S \bullet W) = P(t \text{ produces a true belief} \mid T \bullet S)$, and hence $T \bullet S \bullet W$ is not a statistically

relevant partition of $T \bullet S$. Since by stipulation $T \bullet S \bullet W$ is the only proper partition of $T \bullet S$, $T \bullet S$ is objectively homogenous. And since $T \bullet S$ is the only proper partition of T , $T \bullet S$ is *maximally* objectively homogenous. Therefore, on the Statistical Solution, $T \bullet S$ is the relevant type of t , which is the result we wanted.

The foregoing example illustrates the power of Beebe's Statistical Solution. By requiring that relevant types be objectively homogenous, it avoids the No-Distinction Problem. By requiring that relevant types be maximally objectively homogenous, it avoids the Single Case Problem. And it provides a satisfying response to Conee and Feldman's (1998, p. 4) list of demands: it's a principled solution that doesn't rely on *ad hoc* stipulations, it makes intuitively plausible epistemic classifications, and it's reliabilistic through and through. Moreover, it's rigorous, conceptually sparse, and plausible in its account of cognitive mechanisms. Unfortunately, however, Beebe's solution succumbs to a version of the Single Case Problem and to a related problem raised by Dutant and Olsson (2013).

4. Why Beebe's solution fails

Consider a slightly more fleshed out version of process type T_i : *inferring that one is writing this essay on January 1 at noon (using procedure m_i to solve information-processing problem i_s and executing it on a system of type s_w)*. T_i has one and only one token (t), and t produces a true belief. So $P(t \text{ produces a true belief} \mid t \text{ belongs to } T_i) = 1$; that is, T_i is perfectly reliable. Since T_i has only one token, it doesn't have any proper partitions; thus, T_i is (trivially) objectively homogenous. And since T_i is perfectly reliable, we can plausibly assume that it's *maximally* objectively homogenous, since no type-class (or at least no interesting type-class) of which T_i is a proper subclass is perfectly reliable. So on the Statistical Solution, T_i is the relevant type for t_i . But it seems that T_i is too small to be a genuine relevant type; were T_i the relevant type for t , the belief produced by t would be perfectly justified, which is implausible. The Statistical Solution is therefore susceptible to a version of the Single Case Problem.

Dutant and Olsson (2013, pp. 1354-5) raise a similar problem: call it the *Trivialization Problem*. Consider process type T , and consider the proper subtype of T (call it $T+$) that subsumes all and only those tokens of T that produce true beliefs. Since $T+$ contains *only* those tokens that produce true beliefs, $T+$ is perfectly reliable and thus has no statistically relevant partitions; and since $T+$ contains *all* tokens of T that produce true beliefs, every subclass of T of which $T+$ is a proper partition is less than perfectly reliable, and so $T+$ is maximally objectively homogenous. A parallel argument can be run for the complement of $T+$ (call it $T-$), which is the proper subtype of T that subsumes all and only those tokens of T that produce *false* beliefs. The upshot is this: every maximally objectively homogenous subtype is either perfectly reliable or perfectly unreliable. So for any process token t , if t produces a true belief, its relevant process type is perfectly reliable, and if t produces a false belief, its relevant process type is perfectly unreliable. Hence, “reliability collapses into truth (and anti-reliability into falsehood)” (p. 1364). This result is clearly quite bad for the Statistical Solution.

Dutant and Olsson consider three attempts to salvage the Statistical Solution, but conclude that each attempt either doesn't solve the Trivialization Problem or raises unacceptable new problems. First, they consider whether tightening the notion of statistical relevance will prevent trivialization. Specifically, they consider two ways of tightening statistical relevance: *a*) by restricting it to conditions that increase or decrease the reliability of a process type above or below some threshold (determined contextually or otherwise), and *b*) by restricting it to conditions that “*significantly* [affect] the reliability of a process type” (pp. 1356-7).¹⁹ But as Dutant and Olsson show, *(a)* and *(b)* succumb to a similar trivialization result. If we tighten relevance by way of either *(a)* or *(b)*,

¹⁹ On *(b)*, a type-class should be counted as statistically relevant only if its degree of reliability differs *significantly* from those of the type(s) containing it. The idea is that insignificant differences in reliability—say, the difference between 0.475 and 0.476—do not signal a genuine difference in relevance. The problem with this proposal, aside from Dutant and Olsson's trivialization result, is that there seems to be no principled way to distinguish significant from insignificant differences in reliability. I am indebted to Joshua Myers and Nevin Johnson for raising a similar worry.

then for every type t , if t produces a true belief, t belongs to a reliable process, and if t produces a false belief, t belongs to an unreliable process. This too is undesirable; not every true belief is reliably produced, and not every false belief is unreliably produced. So tightening statistical relevance in either of these ways is of precious little help.

Second, they consider whether one way of “loosening homogeneity” will prevent trivialization. To loosen homogeneity is to impose a condition that keeps subtypes like $T+$ and $T-$ from being “admissible subtypes” (p. 1358). Dutant and Olsson note Wesley Salmon’s (1984) definition of homogeneity, which rules out subtypes that cannot “be ‘determined without reference to’ the outcome we are considering” (Dutant and Olsson, 2013, p. 1358). While specifying homogeneity in this way is supposed to help prevent trivialization, there are issues with Salmon’s proposal. First off, it’s not clear how we can discern, in any given case, whether a subtype can be determined without reference to an outcome. More seriously, the new condition doesn’t rule out certain subtypes whose descriptions have cognitively irrelevant content. Consider Dutant and Olsson’s example:

...even if we conceded that [the process type] “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 (p. 1359).

So even if we’re able to rule out subtypes that refer to a particular outcome, we’re still left with cognitively irrelevant subtypes that are nonetheless statistically relevant. This way of loosening homogeneity, then, won’t do.

Third (and finally), Dutant and Olsson consider another way of loosening homogeneity: by “restrict[ing] admissible types to natural kinds, where a natural kind is thought of as a kind that

appears substantially in some natural law” (p. 1360). This restriction would rule out types like $T+$ and $T-$, since they clearly do not resemble anything like a natural kind. A problem with this proposal is that there seems to be no principled way to determine *which* natural kind is relevant for any given token. What counts as a natural kind varies with different conceptual schemes, and it’s unclear what conceptual schemes have purchase in any given case. Appealing to natural kinds, then, will not suffice.²⁰

The Statistical Solution fails as it currently stands. Dutant and Olsson make a stronger claim, however. They write: “There [is] no quick fix to [the Trivialization Problem]...The conclusion is inevitable: reliabilists have to look elsewhere for a convincing solution to the Generality Problem” (p. 1364).²¹ Dutant and Olsson hold out about as much hope for the Statistical Solution as Conee and Feldman hold out for any solution. But is this attitude warranted? Is there really no (reasonably quick) fix for the Statistical Solution?

5. The New Statistical Solution

The Statistical Solution can be repaired—or so I aim to show. First, I revise condition (c) of the Statistical Answer. Second, I describe the revised condition and show how the resulting solution avoids the pitfalls of the original Statistical Solution. I begin with the revised condition:

New Statistical Answer: For any process token t , T is the relevant process type only if c^* for any class \mathcal{A} of which T is a subclass and in which t falls, T is the broadest *admissible* objectively homogenous subclass of \mathcal{A} , where T is admissible just in case, for any I, M , and

²⁰ For more on natural kinds and the Generality Problem, see Alston (1995), Conee & Feldman (1998, pp. 10-11), and Dutant & Olsson (2013, pp. 1360-62).

²¹ It is worth noting that Olsson (2016) indeed looks elsewhere, namely, to experimental philosophy. While I think that there are significant disadvantages to Olsson’s approach, I do not have the space to discuss them here.

S : if \mathcal{A} is partially defined as $I \bullet M \bullet S$, then the extension of $I_s \bullet M_t \bullet S_u$ is a proper subset of the extension of $I \bullet M \bullet S$.²²

Combining the Tri-Level Answer and the New Statistical Answer yields the *New Statistical Solution*:

New Statistical Solution: For any process token t , T is the relevant process type if and only if *a) t tokens T ; b) for some I_s , M_t , and S_u : T is partially defined as $I_s \bullet M_t \bullet S_u$; and $c^*)$ for any class \mathcal{A} of which T is a subclass and in which t falls, T is the broadest *admissible* objectively homogenous subclass of \mathcal{A} , where T is admissible just in case, for any I , M , and S : if \mathcal{A} is partially defined as $I \bullet M \bullet S$, then the extension of $I_s \bullet M_t \bullet S_u$ is a proper subset of the extension of $I \bullet M \bullet S$.*

The new material in (c^*) can be summed up in a rather ungainly slogan: “No admissibility without a difference in tri-level property graining”. On (c^*) , no subclass T whose tri-level properties are just as general as those of the larger class \mathcal{A} that contains it counts as an admissible subclass of \mathcal{A} . So if t is a token of T , and T has the very same tri-level properties as \mathcal{A} , then T is not the relevant type for t . The intuitive rationale for (c^*) is that, for a type to be relevantly different from some type that properly contains it, the former must differ from the latter with respect to the sort of information-processing problem it solves, the sort of problem-solving procedure it employs, or the sort of system on which the problem-solving procedure in question is executed.

Before moving on to the upshots, let us first get clearer on the notions of extension and graining operative in the New Statistical Solution.²³

Property Extension: The extension of a property X is just the set of objects that instantiate X . So the extension of a property Y is a proper subset of the extension of X just in case the

²² My thanks to an anonymous referee for identifying a problem with an earlier version of (c^*) .

²³ My thanks to an anonymous referee for inviting me to elaborate on my notions of extension and graining.

set of Y -instances is a proper subset of the set of X -instances. *Example:* the extension of “mauve” is a proper subset of the extension of “purple”.

Property Conjunction Extension: The extension of a property conjunction $X \bullet Y$ is the set of objects that instantiate both X and Y . So the extension of a property conjunction $U \bullet V$ is a proper subset of the extension of $X \bullet Y$ just in case the set of $(U \bullet V)$ -instances is a proper subset of the set of $(X \bullet Y)$ -instances. *Example:* the extension of “mauve and floral” is a proper subset of the extension of “purple and floral”.

Tri-Level Property Conjunction Extension: The extension of a conjunction of tri-level properties $I \bullet M \bullet S$ is the set of process types that instantiate I and M and S . So the extension of a tri-level property conjunction $I_s \bullet M_t \bullet S_u$ is a proper subset of the extension of $I \bullet M \bullet S$ just in case the set of process types instancing I_s and M_t and S_u is a proper subset of the set of process types instancing I and M and S . *Example:* compare the most general type class satisfying the description “inferring on the basis of sense perception” with the most general type class satisfying the description “inferring”. The extension of the I -property of the former is a proper subset of the extension of the I -property of the latter, since every solution to an inference-from-sense-perception problem is a solution to an inference problem, but not every solution to an inference problem is a solution to an inference-from-sense-perception problem.²⁴ It follows, then, that the extension of the conjunction $I_s \bullet M_t \bullet S_u$ is a proper subset of the extension of the conjunction $I \bullet M \bullet S$, where the former type class instantiates $I_s \bullet M_t \bullet S_u$ and the latter type class instantiates $I \bullet M \bullet S$.

²⁴ The intuitive gloss provided here is just that: a gloss. I provide a more rigorous account of I -property gaining later in the paper.

Graining: A tri-level property Z_x is *finer grained* than a distinct tri-level property Z just in case the extension of Z_x is a proper subset of the extension of Z . So type class T has finer grained tri-level properties than those of type class A just in case the extension of $I_s \bullet M_t \bullet S_u$ is a proper subset of the extension of $I \bullet M \bullet S$, where T instantiates $I_s \bullet M_t \bullet S_u$ and A instantiates $I \bullet M \bullet S$.

We are now in a position to consider how the New Statistical Solution improves upon Beebe's solution. First, the New Statistical Solution circumvents the Single Case Problem. Again consider T_t : *inferring that one is writing this essay (using procedure m_t to solve information-processing problem i_s and executing it on a system of type s_u) at exactly noon on January 1*. Recall that on the Statistical Solution, T_t is the relevant type for T_t 's only token (t), and hence, the belief produced by t is perfectly justified. But happily, (c^*) avoids this result. On (c^*) , if T_t is a proper subclass of A , and T_t and A have the very same tri-level properties, then for any token of T_t , its relevant type is *not* T_t . And indeed, T_t is in fact a subtype of a larger type that has the same tri-level properties. Compare T_t with T : *process of inferring that one is writing a paper on the Generality Problem (using procedure m_t to solve information-processing problem i_s and executing it on a system of type s_u)*. T 's description is just like that of T_t , but without the cognitively irrelevant details that were nonetheless statistically relevant. And unlike T_t , T has more than one token and isn't perfectly reliable. T is clearly a better candidate for relevance than T_t , and (c^*) reflects that fact. In general, conditions like the date, time and location, say, at which a type is tokened make no difference when it comes to the information-processing problem being solved, the procedure being used, or the type of system on which the procedure is executed. In other words, cognitively irrelevant factors don't bear on tri-level properties. Condition (c^*) rightly rules out these cognitively irrelevant factors.

Second, the New Statistical Solution avoids the Trivialization Problem. To see why, consider another abstract example. Say that type T has five (and only five) tokens: t_1 , t_2 , t_3 , t_4 , and t_5 . The first three tokens (t_1 , t_2 , and t_3) produce true beliefs, and the last two tokens (t_4 and t_5) produce false

beliefs. The problem for the original Statistical Solution was this: T has two maximally objectively homogenous subtypes, one of which is perfectly reliable, and one of which is perfectly unreliable. So on the Statistical Solution, the beliefs produced by t_1 through t_5 are either perfectly justified or perfectly unjustified. But on the New Statistical Solution, if T has the very same tri-level properties as each of its subtypes, then by (c^*) , the relevant type for t_1 through t_5 is T , delivering the desired result: the beliefs produced by t_1 through t_5 are neither perfectly justified nor perfectly unjustified.

I see two interesting objections to the New Statistical Solution, the latter of which deserves particularly close attention. The first objection says that I , M , and S_u can still be assigned definitions so specific that only one token falls under the instantiating type, thereby raising the Single Case Problem all over again. My initial response is this: *i*) an admissible description of a tri-level property can include only cognitively relevant factors; *ii*) in general, descriptions of single-token types include cognitively irrelevant factors (e.g., dates and times); *iii*) so, condition (c^*) should deliver the verdict that all (or virtually all) single-token types are irrelevant.

The objector may retort, however, that I , M , and S_u need not have cognitively irrelevant factors built into their descriptions in order to have only one token in their extensions. Rather, the information-processing problem, the procedure, and the system type might be defined so finely that only one token *in fact* falls under them. We can, for example, revise T_1 so that it has only one token and includes no cognitively irrelevant factors in its description. Consider our new T_1 (call it T_1^*): *the process of inferring that one is writing a paper on the Generality Problem, using a procedure exactly similar to that which the author of this paper uses to solve an exactly similar problem and executing it on a system that's exactly like this author's information-processing system.* T_1^* has only one token; and yet, since it doesn't have any cognitively irrelevant factors built into its description, T_1^* still comes out admissibly, maximally, and objectively homogenous. So the New Statistical Solution still succumbs to the Single Case Problem.

The above objection is *prima facie* plausible, but ultimately unsuccessful. We should reject the assumption that T_i^* has only one token. To say that T_i^* has only one token is to say that two individuals cannot solve exactly similar information-processing problems using exactly similar procedures and executing them on exactly similar system types. But why suppose this? Why think that distinct individuals can't token the same maximally specific type defined only in terms of tri-level properties? We have no good reason to think this; indeed, it appears that being multiply realizable is just part of what it is to be a relevant process type. Consider Beebe's rationale for excluding implementation-level properties from the stock of tri-level properties:

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 ... *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 cognitive relevant facts is precisely why cognitive scientists look to higher-order functional descriptions that abstract from many of the physical details of cognitive systems (pp. 184-5, emphasis added).

The upshot is that computational and algorithmic properties are, by their very nature, shareable. Indeed, that's part of what makes the types that realize these properties viable candidates for relevance. Implementation-level properties, on the other hand, are not sharable, which is what makes the physically characterized process types specified by Conee and Feldman (1998) unviable candidates for relevance. The lesson here is that the New Statistical Solution has a built-in check against the Single Case Problem, for it entails that relevant process types are multiply realizable. This built-in check is neither *ad hoc* nor question begging; it follows, rather, from the plausible assumption

that the physical properties of an individual cognitive system are irrelevant when it comes to picking out relevant process types.²⁵

The second objection says that, unlike Beebe's Statistical Solution, the New Statistical Solution cannot account for the role that environments play in determining type reliability.²⁶ To set the stage, let us return to some example types that were discussed in section 3: inferring (T), inferring on the basis of sense perception under favorable conditions ($T\bullet S$), and inferring on the basis of sense perception under favorable conditions on a Wednesday ($T\bullet S\bullet W$). A virtue of both Beebe's Statistical Solution and the New Statistical Solution is that they pick out $T\bullet S$ as the relevant type for token t of $T\bullet S\bullet W$. On Beebe's Statistical Solution, $T\bullet S$ is the relevant type since $T\bullet S$ is the broadest objectively homogenous subclass of T that includes t . On the New Statistical Solution, $T\bullet S$ is the relevant type since $T\bullet S$ is both the broadest objectively homogenous subclass of T that includes t and an admissible subtype of T .

But let us replace T with another type of which $T\bullet S$ and $T\bullet S\bullet W$ are subtypes—*inferring on the basis of sense perception* (T^*)—and ask whether both solutions still deliver the correct verdict, namely, that $T\bullet S$ is the relevant type. Because $T\bullet S$ is a statistically relevant partition of T^* , Beebe's Statistical Solution still picks out $T\bullet S$ as relevant. It seems, however, that the New Statistical Solution picks out T^* as relevant. $T\bullet S$, it appears, does not have finer grained tri-level properties than those of T^* ; $T\bullet S$ neither solves a narrower processing problem nor employs a narrower problem-solving procedure nor executes its problem-solving procedure on a narrower type of system than does T^* . More precisely, if $T\bullet S$ is partially defined as $I_s\bullet M_t\bullet S_{\bullet}$, and T^* is partially defined as $I\bullet M\bullet S$, then the

²⁵ A persistent objector, acquiescing the multiple realizability of relevant process types, might still complain that, while T_i^* can be realized more than once, it probably won't be realized more than once, given how specific T_i^* is. All that's needed for the Single Case Problem to remain is for T_i^* to have only one token *in fact*. So according to the objector, the New Statistical Solution is still susceptible to the Single Case Problem. My response is that, if this is indeed a genuine problem, it is a problem only for probabilistic or frequentist reliabilism, not for the New Statistical Solution *per se*. Modal reliabilism, for example, has resources to countenance the objector's complaint and is fully compatible with the New Statistical Solution. For more on modal reliabilism, see Becker (2007, pp. 21-39) and Baumann (2009).

²⁶ My thanks to two anonymous referees for raising this objection.

extension of $I_s \bullet M_i \bullet S_n$ is not a proper subset of the extension of $I \bullet M \bullet S$; that is, $T \bullet S$ is not an admissible subtype of T^* . But if $T \bullet S$ is not an admissible subtype of T^* , the New Statistical Solution picks out T^* as the relevant type. This, of course, is an untoward result for the defender of the New Statistical Solution.

The source of the apparent problem is this: “under favorable conditions” is an external condition that bears on the reliability of $T \bullet S$. But tri-level conditions are *internal* conditions. To impose, then, an admissibility condition like (c^*) is to deny the role that external conditions play in determining type relevance. It should be clear, to any reliabilist at least, that external conditions are, in a word, relevant to relevance; whether an act of sense perception occurs under favorable conditions, for example, clearly bears on the reliability of the produced inference. The New Statistical Solution, however, does not account for the role external conditions play in determining reliability, nor does it honor the externalist spirit of reliabilism. Or so the objection goes.

If the New Statistical Solution indeed had no place for external conditions in determining relevance, I too would be resistant to the view. But the New Statistical Solution does have a place for external conditions. To see how, let us begin by considering a recent remark from a leading proponent of the Tri-Level Hypothesis, cognitive scientist Michael Dawson:

At the computational level, one uses a formal vocabulary to provide a rigorous description of *input-output mappings*. At the algorithmic level, a procedural or behavioural vocabulary is employed to describe the algorithm being used to calculate a particular input-output mapping (2013, p. 48, emphasis added).

This terse excerpt suggests what we might call an *externalist* conception of computational analysis and an *internalist* conception of algorithmic analysis. Computational analysis is external because it “reaches outside” to the world beyond the cognitive system. Algorithmic analysis is internal because it “stays in-house”. Allow me to elaborate on this crude distinction.

Recall that on Beebe's Tri-Level Definition, I_i is the property of solving information problem i , and M_m is the property of using procedure m , to solve an information-processing problem. With Dawson's treatment in mind, we can roughly understand I -properties in terms of input-output mappings and M -properties in terms of functions; thus, an I -property is the property of having an input-output mapping of a certain sort, and an M -property is the property of applying a certain sort of function to the input(s) in question. The significance of this analysis lies in the way *inputs* figure into I -properties. What makes an I -property this or that sort of I -property is the relation that exists between input and output. Moreover, what inputs a system has often (if not always) causally depends on the system's environment. Therefore, what I -properties a system has isn't simply an internal matter; a system's I -properties change in response to varying environmental conditions that present diverse inputs.²⁷ Indeed, this is one of the features that distinguishes I -properties from M -properties; while the latter are merely functional and involve nothing "outside the system", the former map inputs to outputs and thus "reach beyond the system" to the surrounding environment.

Herein lays the key to resolving the second objection. Since I -properties are to be understood in terms of input-output mappings, I -properties are partly to be individuated by inputs. So contrary to the second objection, the types corresponding to the descriptions "inferring on the basis of sense perception (T^*)" and "inferring on the basis of sense perception under favorable conditions" ($T \bullet S$) do *not* have the very same tri-level properties, since the set of inputs associated with T^* is distinct from the set of inputs associated with $T \bullet S$; and so, T^* and $T \bullet S$ have different I -properties. Importantly, the difference between T^* 's and $T \bullet S$'s I -properties is a difference in *graining*. $T \bullet S$'s I -property is finer grained than T^* 's I -property, since the set of inputs associated with $T \bullet S$ is a proper subset of the set of inputs associated with T^* , and since I -property graining is partly a

²⁷ My thanks to an anonymous referee for drawing my attention to the relationship between sensory inputs, environmental conditions, and internal processing.

function of (what we might call) *I*-input graining. $T \bullet S$ is therefore an admissible subtype of T^* on the New Statistical Solution.

A brief word is due to those dubious of the claim that *I*-properties can differ with respect to graining.²⁸ We have seen that *I*-properties can be distinguished by their inputs. But if the only difference between two *I*-properties—call them I_1 and I_2 —is a difference in input, and if I_2 is finer grained than I_1 , the implication seems to be that the input associated with I_2 is finer grained than the input associated with I_1 . But does it make sense to speak of one input as being finer grained than another? Aren't inputs just different, period?

I do think it appropriate to analyze inputs in terms of graining, and this for four reasons. First, in the perception literature, perceptual experiences are commonly analyzed in terms of graining.²⁹ So far from being an unwelcome artifact of an *ad-hoc* maneuver, comparing inputs in terms of graining is quite natural, at least within the paradigm of contemporary philosophy of perception. Second, we can legitimately analyze sensory inputs in terms of phenomenal properties.³⁰ If, then, inputs are analyzable in terms of *properties*, they are also analyzable in terms of *graining*, *per*

²⁸ My thanks to an anonymous referee for inviting me to say more about this.

²⁹ See Kelly (2001), Neale (1998), Siegal (2016), and Tye (1995). There is a long-standing debate (which Siegal [2016] deftly summarizes) over whether perceptual inputs are maximally fine-grained. At first, this might seem to spell trouble for my account; for if perceptual inputs are maximally fine-grained, then it would seem that *I*-properties are also maximally fine-grained, and hence, that tri-level property conjunctions are maximally fine-grained. But doesn't this raise the Single Case Problem all over again? The answer is *no*, for two reasons. First, admissibility is a necessary but insufficient condition for relevance. That a token instances a maximally specific type does not imply that the maximally specific type is the relevant type for the token in question. Second, even if perceptions are maximally fine-grained, we can modify the account so that *I*-properties are not systematically maximally fine-grained. The debate over perceptual graining can be framed roughly as a question: at what stage in the perceptual process are concepts applied to raw sensory data? Regardless of our answer, we can simply stipulate that the inputs associated with *I*-properties are *post-concept-application inputs*. So say, for example, that Smith and Jones are looking at the same shade of mauve. Say further that Smith, but not Jones, can distinguish mauve from purple; Smith has a concept of mauve that Jones lacks. For the purposes of the New Statistical Solution, we can say that Smith is being appeared to mauvely and Jones is being appeared to purplely; that Smith and Jones are experiencing different, non-maximally fine-grained inputs; and hence, that Smith's and Jones's perceptual belief-forming processes token different, non-maximally specific process types. Construed in this way, inputs need be no finer grained than concepts applied to raw sensory data.

³⁰ In his psychologically-informed discussion of phenomenal properties, Clark (2008) defines a phenomenal property as “very roughly, a characteristic of sensible appearance: a quality that qualifies how things *appear*” (p. 407). This is the sense of “phenomenal property” that I have in mind, though I take the New Statistical Solution to be compatible with a wide variety of accounts of phenomenal properties. The question of how fine phenomenal properties are needn't concern us here. (See footnote 29).

our operative definition of “graining”. Third, the notion that one perception can be finer grained than another is, I think, fairly intuitive. That “clear visual perception” should come out finer grained than “visual perception” is unsurprising; and happily, this is just what the New Statistical Solution suggests. Fourth, analyzing inputs in terms of graining makes for a nice isomorphism between perceptions and objects of perceptions. Just as “mauve” is finer grained than “purple”, so being appeared to mauvely is finer grained than being appeared to purplely. It is thus natural to treat inputs in a way that mirrors the external world.

Let us put the foregoing analysis to work with a final example. Consider two agents: Smith and Jones. Smith perceives a duck under favorable conditions and correctly infers that there is a duck in the pond. Jones perceives a duck under unfavorable conditions, but nonetheless correctly infers that there is a duck in the pond. Smith’s and Jones’s inferences both token T^* , but only Smith’s inference tokens $T \bullet S$. Since Smith’s belief is better justified than Jones’s belief (or so reliabilism would have it), the New Statistical Solution had better determine that $T \bullet S$ is the relevant type for Smith; otherwise, the New Statistical Solution runs afoul of the No-Distinction Problem. Happily, the New Statistical Solution delivers the correct result. The reason that Smith’s belief is better justified than Jones’s is at least partly because Smith had, as it were, a better look at the duck. Perhaps it was a sunny day for Smith and a foggy day for Jones, or perhaps Smith had her glasses on and Jones did not. What is important is that Smith and Jones experienced different visual inputs, despite having generated the same output inference. If, then, Smith and Jones experienced different inputs, and I -properties are at least partly differentiated by inputs, then Smith’s and Jones’s respective process tokens instance process types with different I -properties; and so, by (i^*) , the relevant process types for Smith and Jones are different. This is just the result we wanted, for it correctly implies that Smith’s inference is better justified than Jones’s inference. In this way, the New

Statistical Solution is rightly sensitive to environmental conditions and how they bear on type reliability.

6. Concluding remarks

I have advanced a new solution to the Generality Problem—one that builds upon Beebe’s Statistical Solution and countenances Dutant and Olsson’s criticisms. The New Statistical Solution retains the virtues and avoids the problems of Beebe’s account. Moreover, the New Statistical Solution neither significantly complicates the Statistical Solution nor increases its proponent’s theoretical commitments; it merely adds a plausible and theoretically motivated admissibility condition. But even if I have failed to show, once and for all, that the Generality Problem is solved, I trust that I have at least shown this: statistical approaches to the Generality Problem are fit not for abandonment, but for further development.³¹

References

- Alston, W. 1995. ‘How to Think About Reliability’. *Philosophical Topics*, 23.1: 1-29.
- Baumann, P. 2009. ‘Reliabilism—Modal, Probabilistic or Contextualist’. *Grazer Philosophische Studien*, 79: 77-89.
- Becker, K. 2007. *Epistemology Modalized*. New York: Routledge.
- . 2008. ‘Epistemic Luck and the Generality Problem’. *Philosophical Studies*, 139.3: 353-66.
- Beebe, J. R. 2004. ‘The Generality Problem, Statistical Relevance and the ‘Tri-Level Hypothesis’. *Noûs*, 38.1: 177-95.
- Bishop, M. 2010. ‘Why the Generality Problem is Everybody’s Problem’. *Philosophical Studies*, 151.2: 285-98.
- Clark, A. 2008. ‘Phenomenal Properties: Some Models from Psychology and Philosophy’. *Philosophical Issues*, 18: 406-25.
- Comesaña, J. 2006. ‘A Well-Founded Solution to the Generality Problem’. *Philosophical Studies*, 129.1: 27-47.
- Conee, E. 2013. ‘The Specificity of the Generality Problem’. *Philosophical Studies*, 163.3: 751-62.
- Conee, E. and Feldman, R. 1998. ‘The Generality Problem for Reliabilism’. *Philosophical Studies*, 89.1: 1-29.

³¹ My thanks to Nathan Ballantyne, James Beebe, Christopher Buford, Robert Duffy, Alvin Goldman, Julia Hass, Quinn Harr, Ben Henke, Eric LaRock, Kevin Richardson, Coran Stewart, and anonymous reviewers for helpful comments on earlier drafts.

- . 2001. 'Internalism Defended'. In H. Kornblith (ed), *Epistemology: Internalism and Externalism*, 231-60. Malden, MA: Blackwell.
- Dawson, M. R. W. 1998. *Understanding Cognitive Science*, Oxford: Blackwell.
- . 2013. *Mind, Body, World: Foundations of Cognitive Science*, Edmonton: Athabasca University Press.
- Dutant, J. and Olsson, E. J. 2013. 'Is There a Statistical Solution to the Generality Problem?' *Erkenntnis*, 78.6: 1347-65.
- Feldman, R. 1985. 'Reliability and Justification'. *The Monist*, 68: 159-74.
- Goldman, A. I. 1979. 'What is Justified Belief?' In *Reliabilism and Contemporary Epistemology*, 29-49. Oxford: Oxford University Press.
- . 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- . 2011. 'Reliabilism (Spring 2011 Edition)'. In *Reliabilism and Contemporary Epistemology*, 68-94. Oxford: Oxford University Press.
- . 2012. 'Introduction'. In *Reliabilism and Contemporary Epistemology*, 3-28. Oxford: Oxford University Press.
- . 2015. 'Reliabilist Epistemology'. In *Stanford Encyclopedia of Philosophy*, <<http://plato.stanford.edu/entries/reliabilism/>>.
- Kelly, S. D. 2001. 'The Non-Conceptual Content of Perceptual Experience: Situation Dependence and Fineness of Grain'. *Philosophy and Phenomenological Research*, 62: 601-8.
- Marr, D. 1982. *Vision*. San Francisco: W.J. Freeman.
- Matheson, J. D. 2015. 'Is There a Well-Founded Solution to the Generality Problem?' *Philosophical Studies*, 172.2: 459-68.
- McEvoy, M. 2005. 'Belief-Independent Processes and the Generality Problem for Reliabilism'. *Dialectica*, 59.1: 19-35.
- Neale, S. 1998. 'Grain and Content'. *Philosophical Issues*, 9: 353-8.
- Olsson, E. J. 2016. 'A Naturalistic Approach to the Generality Problem'. In B. McLaughlin and H. Kornblith (eds), *Goldman and His Critics*, 178-99. Malden, MA: Blackwell.
- Prinz, J. 2007. 'The Intermediate Level Theory of Consciousness'. In S. Schneider and M. Velmans (eds), *Blackwell Companion to Consciousness*, 247-60. Oxford: Blackwell.
- . 2012. *The Conscious Brain*. Oxford: Oxford University Press.
- Salmon, W. C. 1984. 'Statistical Explanation'. In W.C. Salmon, R. Jeffrey, and J. Greeno (eds), *Statistical Explanation and Statistical Relevance*, 29-88. Pittsburgh: University of Pittsburgh Press.
- Schmitt, F. F. 1992. *Knowledge and Belief*. New York: Routledge, Chapman, and Hall.
- Siegel, S. 2016. 'The Contents of Perception'. In *Stanford Encyclopedia of Philosophy*, <<http://plato.stanford.edu/entries/perception-contents/>>.
- Tye, M. 1995. 'Non-Conceptual Content, Richness, and Fineness of Grain'. In T. Gendler and J. Hawthorne (eds) *Perceptual Experience*, 504-26. Oxford: Oxford University Press.
- Wunderlich, M. E. 2003. 'Vector Reliability: A New Approach to Epistemic Justification'. *Synthese*, 136.2: 237-62.