

# Rejecting the New Statistical Solution to the Generality Problem

Jeffrey Tolly

\*Penultimate Draft\*

Forthcoming in *Episteme*. Please only cite published version. Official pre-published version available here:

DOI:

<https://doi.org/10.1017/epi.2019.18>

The generality problem is one of the most pressing challenges for process reliabilism about justification. Thus far, one of the more promising responses is James Beebe's tri-level statistical solution. Despite the initial plausibility of Beebe's approach, the tri-level statistical solution has been shown to generate implausible justification verdicts on a variety of cases. Recently, Samuel Kampa has offered a *new statistical solution* to the generality problem. Kampa argues that the new statistical solution overcomes the challenges that undermined Beebe's original statistical solution. However, there's good reason to believe that Kampa is mistaken. In this paper, I show that Kampa's new statistical solution faces problems that are no less serious than the original objections to Beebe's solution. Depending on how we interpret Kampa's proposal, the new statistical solution *either* types belief-forming processes far too narrowly, *or* the new statistical solution fails to clarify the epistemic implications of reliabilism altogether. Either way, the new statistical solution fails to make substantive progress towards solving the generality problem.

Generality Problem, Process Reliabilism; Justification; Knowledge

## 1. Introduction

The generality problem is one of the most important challenges to process reliabilism about epistemic justification.<sup>1</sup> In brief, the generality problem poses the following question: out of all the process types exemplified by a given process token, *which* types are epistemically relevant for determining justification? Presumably, without an answer to this question, it's very hard

---

<sup>1</sup> Early on, Alvin Goldman described the generality problem as a "critical problem" for the reliabilist analysis of justification (1979). More recently, after surveying the literature on reliabilism throughout the past four decades, Beddor and Goldman identify the generality problem as one of the top six "problems," or, "objections" to reliabilism (2015).

to tell what epistemic implications (if any) reliabilism will have on particular cases of belief formation.<sup>2</sup> Hence, any satisfactory solution to the generality problem must offer an informative theory of process type relevance.

As the past four decades have demonstrated, substantive progress on the generality problem has been remarkably hard to come by.<sup>3</sup> That said, some responses to the generality problem have garnered more attention—and appear more promising—than others. One of these responses is James Beebe’s tri-level statistical solution to the generality problem. However, despite the initial plausibility of Beebe’s approach, Julien Dutant and Erik Olsson have shown that the tri-level statistical solution entails intuitively implausible justification verdicts on a variety of cases (2013).

Samuel Kampa has recently offered a new proposal for repairing Beebe’s original solution (2018). Kampa calls it the *new statistical solution* to the generality problem. After presenting the new statistical solution, Kampa argues that it successfully overcomes the challenges that undermined Beebe’s original statistical solution. However, there’s good reason to believe that Kampa is mistaken. In this paper, I argue that Kampa’s new statistical solution fails to make substantive progress towards solving the generality problem.

In §2, I present Beebe’s statistical solution in more detail. After this, I explain Dutant and Olsson’s contention that the types identified as relevant according to Beebe’s theory are far too descriptively *narrow* to be the correct relevant types. In §3, I present Kampa’s new statistical solution to the generality problem and discuss Kampa’s explanation for why his new solution avoids the problems that plague Beebe’s theory. In §4, I offer my main criticism of Kampa’s new statistical solution. In particular, I show that the new statistical solution faces a dilemma. If we interpret the elements of the new statistical solution in an unqualified manner, then the new statistical solution falls prey to a straightforward counter example by countenancing types that are far too narrow. On the other hand, if we interpret the new statistical solution in a qualified manner, then the theory in its current form offers us, at best, scant insight into the nature of process type relevance. Either way, Kampa’s theory fails as a genuine solution to the generality problem.

## 2. The original tri-level statistical solution

### 2.1 The tri-level condition and Beebe’s statistical solution

James Beebe invokes key notions from cognitive science to formulate his answer to the generality problem. According to neuro-scientist David Marr, cognitive processes can be analyzed at three levels of description: The information problem (I) being solved in the process, the method (M) used to solve that problem, and the cognitive system (S) used to

---

<sup>2</sup> See Richard Feldman (1985: 165) for further discussion on the burden reliabilists have to supply a theory of type relevance.

<sup>3</sup> For example, see Conee and Feldman (1998) for lengthy treatment on the failures of numerous extant responses to the generality problem. In addition, see Adler and Levin (2002) and Comesaña (2006) for newer responses to the generality problem. See Conee and Feldman (2002) for a criticism of Adler and Levin’s proposal, and see Matheson (2015) for a criticism of Comesaña’s proposal.

execute that method.<sup>4</sup> As Beebe clarifies, the the cognitive method (M) is the *algorithm* used to solve the information problem, and the cognitive system (S) that solves the information problem is the *cognitive architecture* that executes the algorithm (Beebe 2004: 182).<sup>5</sup> Beebe contends that the (I), (M), and (S) properties of belief-forming process tokens are epistemically relevant features that determine (at least partially) whether a given process token generates a justified belief (2004: 180). Beebe incorporates this idea by positing the *tri-level condition* for process type relevance.

The tri-level condition:

The reliability of a cognitive process type T determines the justification of any belief token produced by a cognitive process token t that falls under T only if all of the members of T:

- a) solve the same type of information-processing problem *i* solved by t;
- b) use the same information-processing procedure or algorithm *t* used in solving *i*; and
- c) share the same cognitive architecture as t.<sup>6</sup>

Some brief clarifications are in order. In Beebe’s terminology, the belief-forming process tokens that indeed token (i.e., instantiate) some process type just are the tokens that “fall under” that type. Beebe describes the class of tokens falling under some type T as the “members of T.” This convention is reasonable enough, as one can straightforwardly view types as corresponding to a particular extension, where in this case the extension of a process type is the *class* of process tokens that instantiate that type.<sup>7</sup>

The tri-level condition posits the following constraint on a token t’s relevant type T: all of the tokens that fall under T must have the same (I), (M), and (S) properties that t has. Importantly, for Beebe this is a *partial* definition of token’s relevant type, because it only presents a necessary (but not sufficient) condition on being a member of T’s extension.

In addition to the tri-level condition, Beebe saw that he’d need to place further restrictions on relevant types so as to avoid what Richard Feldman calls the “no distinctions problem.”<sup>8</sup> According to Feldman, a given theory of type relevance succumbs to the no-distinctions problem when the types it identifies as relevant are too broad. Feldman notes that responses

---

<sup>4</sup> See Marr (1982). Also, see Beebe (2004: 181-3) for further explanation on the tri-level approach to analyzing cognitive processes.

<sup>5</sup> Beebe draws this terminology from Dawson (1998).

<sup>6</sup> Beebe (180). As Beebe clarifies, the cognitive architecture (S) property that partially defines a relevant type should not be understood as a *physical* description of the system that implements/executes an information-processing problem (185). Taking on-board the plausible assumption that relevant process types are *multiphysically-realizable* by different physical structures, Beebe insists that the (S) properties mentioned in the tri-level condition are “higher-order *functional* descriptions that abstract from many of the *physical* details of cognitive systems” (185, emphasis mine). Seeing as how Kampa’s response to the generality problem “builds upon” Beebe’s response, I will assume that Kampa adopts the same functional, non-physical construal of (S) properties (Kampa 2018: 229).

<sup>7</sup> It’s important to note that, in some places, Beebe simply identifies types with classes of tokens (for example, see the quote below from 187-8). For our purposes, not much hangs on whether types are identical to classes of tokens or merely *have* classes of tokens as extensions.

to the generality problem that suffer from the ‘no-distinctions problem’ end up entailing the same degree of justification for various beliefs that intuitively should have *different* degrees of justification (Feldman 1985: 161). For instance, consider the following three process types:

- [1] visual belief formation
- [2] visual belief formation under good lighting conditions
- [3] visual belief formation under bad lighting conditions

*Ceteris Paribus*, tokens that instantiate [2] produce beliefs with a greater degree of justification than tokens that instantiate [3]. A reliabilist would explain this fact by noting that [2] is more reliable than [3]. However, a theory of type relevance that identifies broad types like [1] as being the relevant type for *any* token instance of visual belief formation will entail the implausible result that beliefs produced by vision under good lighting conditions will have the same degree of justification as beliefs produced by vision under poor lighting conditions. Hence, we should reject such a theory of type relevance that countenances relevant types that are too broad.

In order to avoid the no-distinctions problem, Beebe adds an additional *statistical* constraint to his theory of type relevance:

Let A be the broadest process type that satisfies the tri-level condition for some process token t...I argue that the relevant process type for some t is the subclass of A which is the broadest objectively homogeneous subclass of A within which t falls. A subclass S is objectively homogeneous if there are no statistically relevant partitions of S that can be effected (2004: 187-8).

Let’s begin by examining the key concepts of this statistical constraint. First, recall that Beebe refers to process types as classes of process tokens. Classes can have proper sub-classes within them, and Beebe uses the notion of a class “partition” to refer to a proper-subclass of some broader class. In this way, a narrower type like [2] is an example of a partition of a broader type like [1].

To understand objective homogeneity, we must first understand the notion of a “statistically relevant” partition. Let A represent a broad process type, and let S represent a proper sub-class type of A. S is a statistically relevant partition of A if and only if S’s degree of reliability differs from A’s degree of reliability (Beebe 188). According to Beebe, a type’s degree of reliability is simply the “probability” that a token generates a true belief given that it instantiates that type (188). For example, given that [2] is a partition of [1] and that (*ceteris paribus*) [2] is more reliable than [1], it follows that [2] is a statistically relevant partition of [1].

Given this notion of a statistically relevant partition, we can understand objective homogeneity as follows: a given type T is objectively homogeneous just if T has no proper sub-types with degrees of reliability that differ from T’s degree of reliability (189).

With both the tri-level condition and the statistical constraint in place, we can now state Beebe’s tri-level statistical solution to the generality problem (TS).

- TS      T is the relevant type for a given token t if and only if
- (a) t tokens T
  - (b) T satisfies the tri-level condition
  - (c) T is the broadest objectively homogenous subclass of the broadest type satisfying the tri-level condition (relative to t)

According to Beebe, the statistical constraint (c) makes the types identified as relevant narrow enough so that TS avoids Feldman’s no-distinctions problem.

## 2.2 Problems with Beebe’s Tri-Level Statistical Solution

Despite its initial plausibility, there are serious problems with TS. Julien Dutant and Erik Olsson present what they call the “trivialization problem” against TS (2013: 1354-5). Consider any type T of a given token t, where T satisfies the tri-level condition. Now, consider the proper subclass of T denoted by T+, where T+ is comprised of all and only the tokens of T that produce *true* beliefs. T+ is perfectly reliable, so not only is T+ statistically relevant, but it also won’t contain any partitions with different degrees of reliability. Moreover, if we suppose that t itself generates a true belief, then t instantiates T+. In this case, TS entails that T+ is the relevant type for t. But given that T+ has a maximal (100%) degree of reliability, it follows that the belief produced by t will be perfectly justified. This schematic description of TS’s workings highlights how *any* true belief, according to TS, will automatically (and *trivially*) be perfectly justified.

This is an implausible result. For instance, TS would entail that any true belief formed on the basis of a coin flip would be perfectly justified. In essence, Dutant and Olsson’s objection shows how TS avoids the no-distinctions problem at the cost of falling prey to the opposite worry for theories of type relevance: what Feldman calls “the single-case problem” (1985: 161). A theory of type relevance suffers from the single-case problem insofar as it countenances types (as being relevant) that are too narrow. According to theories of type relevance that fall prey to the single-case problem, virtually any token that generates a true belief will have a relevant type with maximal (or near-maximal) reliability, and any token that generates false belief will have a relevant type that has maximal (or near-maximal) *unreliability* (161).

Dutant and Olsson consider various strategies for either tightening restrictions on statistical relevance, or loosening restrictions on homogeneity so as to salvage TS in some form or fashion.<sup>9</sup> Ultimately, Dutant and Olsson show there to be serious flaws with all of these repair strategies. While I lack the space to discuss these strategies here, it’s important to note that Kampa accepts the failure of these repair proposals. This leads Kampa to present his own unique strategy for repairing the tri-level statistical solution to the generality problem.

---

<sup>9</sup> (Beebe 2004: 190)

### 3. Kampa's New Statistical Solution

Kampa's main approach for avoiding the problems raised by Olsson and Dutant involves adding a further constraint on the broadest objectively homogenous subclass that can constitute a token's relevant type. Kampa calls this proposal the *New Statistical Solution* (NS).

- NS     For any process token  $t$ ,  $T$  is the relevant process type for  $t$  if and only if
- a)  $t$  tokens  $T$
  - b)  $T$  satisfies the tri-level condition
  - c)  $T$  is the broadest objectively homogenous *admissible* subclass of the broadest type satisfying the tri-level condition under which  $t$  falls (236)

Kampa defines *admissibility* as follows:

#### Admissibility

Where  $A$  is a type partially defined by tri-level properties  $[I_A, M_A, S_A]$  and  $T$  is a type partially defined by tri-level properties  $[I_B, M_B, S_B]$ ,  $T$  is an admissible subclass of  $A$  just in case the extension of  $[I_B, M_B, S_B]$  is a *proper* subclass of the extension of  $[I_A, M_A, S_A]$ .<sup>10</sup>

Kampa claims that “[t]he new material in [condition (c) of NS] can be summed up in a rather ungainly slogan: ‘No admissibility without a difference in tri-level property graining’” (236). According to the admissibility constraint on relevant types, in order for  $T$  to be the relevant type for token  $t$ , it cannot be the case that  $T$  is partially defined by the exact same tri-level properties that partially define some broader type  $A$  (instantiated by  $t$  as well) of which  $T$  is a proper sub-type. In order for a type  $T$  to be admissible,  $T$  must be partially defined by a tri-level property ( $I$ ,  $M$ , or  $S$ ) that is finer-grained than a tri-level property that partially defines type  $A$ . According to Kampa, a given tri-level property  $F_1$  is finer-grained than a distinct tri-level property  $F_2$  if and only if the extension of  $F_1$  is a proper-subset of the extension of  $F_2$  (237).

NS successfully avoids Dutant and Olsson's trivialization problem. This is because it's possible for tokens with true-belief outputs and tokens with false-belief outputs to share all of the same exact tri-level properties. So, while it might be the case that a given token instantiates *some* type that's comprised of all and only true-belief producing instances, there's no guarantee that this type will be admissible. Hence, it's not the case that types like  $T^+$  will automatically count as relevant according to NS (238). Therefore, NS doesn't entail that any true belief will be maximally justified.

According to Kampa, the fact that tri-level properties can be more or less fine-grained allows NS to avoid the no-distinctions problem. While Kampa doesn't specifically discuss how  $M$  or  $S$ -properties can be more or less fine-grained, he goes to some lengths to argue that  $I$ -property graining is both coherent and relatively straightforward. Furthermore, he argues that  $I$ -property graining plays a key role in explaining how NS delivers intuitively correct

---

<sup>10</sup> (Kampa 2018: 234-6) This definition is a paraphrase taken from Kampa's formulation of what he calls the “New Statistical Answer” (236).

justification verdicts on particular cases of belief formation. Following the work of Michael Dawson, Kampa suggests that we define information problems in terms of “input-output” mappings (Kampa 2018: 240, Dawson 2013:48). In this way, “I-properties can be distinguished by their inputs” (241). Let’s say that I-property  $I_a$  has an input-output mapping that is a proper subset of another I-property  $I_b$ ’s input-output mapping. In this case,  $I_a$  is finer-grained than  $I_b$ .

Importantly, Kampa does *not* give a detailed definition or informative analysis that outlines all of the possible ways in which input-properties individuate I-properties and thus determine I-property graining. However, he does give two distinct examples of I-property graining: graining with respect to the phenomenal properties of inputs, and graining with respect to the environmental properties of inputs (241-2). I will quickly address both examples in turn.

Kampa gives the following description of how inputs to perceptual belief-forming processes—perceptual experiences—can be more or less fine-grained according to their *phenomenal* properties:

[W]e can legitimately analyze sensory inputs in terms of phenomenal properties. If, then, inputs are analyzable in terms of properties, they are also analyzable in terms of graining, per our operative definition of “graining”...[T]he 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...[A]nalyzing 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. (242)

Given how NS determines relevant types on this basis of phenomenal-property input graining, it would seem that NS can straightforwardly account for the fact that, *ceteris paribus*, beliefs produced by clear visual perception are more justified than beliefs produced by mere visual perception. Plausibly, the class of visual belief-forming process tokens featuring clear, non-blurry phenomenology has a higher truth-ratio than the broader class of all visual belief-forming tokens. This would make the type [clear visual perception] a statistically relevant subclass of [visual perception]. And, since phenomenal property graining can ground I-property graining, the sub-class of [clear visual perception] would count as admissible according to NS.

Kampa also states that I-properties can be more or less fine-grained with respect to the *environmental* features of input experiences. In applying this idea, Kampa asks us to consider the following two types, which I’ll denote as  $T^*$  and  $T \bullet S$  for short:

$T^*$             [inferring on the basis of sense perception]  
 $T \bullet S$         [inferring on the basis of sense perception under favorable conditions]<sup>11</sup>

Given that *being formed under favorable environmental conditions* is, plausibly, a justification-determining feature of a belief, Kampa notes that  $T \bullet S$  is intuitively the correct relevant type

---

<sup>11</sup> See Kampa (2018: 239-40).

for process tokens that instantiate  $T \bullet S$  (240). Kampa argues that NS can deliver this result in virtue of how I-properties can be more or less fine-grained due to the external/environmental properties of experiential inputs.

[W]hat 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...[I-properties] map inputs to outputs and thus "reach beyond the system" to the surrounding environment. (240-1)

According to this account of I-properties, it's possible to individuate an I-property on the basis of an input-output mapping whose input perceptual experiences were all *formed under a specific environmental condition C*. For Kampa's example, let C denote *favorable environmental conditions for sense perception*. This input-output mapping would constitute an I-property that 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  $T^*$ " (241). Kampa notes that this fact makes  $T \bullet S$  "an admissible subtype of  $T^*$  on the New Statistical Solution," thus allowing NS to pick out  $T \bullet S$  as a relevant type for tokens that instantiate it (241).

In sum, Kampa defends the idea that I-properties can be more or less fine-grained on the basis of either the phenomenal features of inputs or the environmental features of inputs. By his lights, this allows NS to countenance admissible types that allow his solution to the generality problem to avoid the no-distinctions problem.

## 4. Problems with the New Statistical Solution

For whatever merits NS has, I contend that it fails as an informative account of type relevance. To begin, it's unclear exactly how we are to interpret NS. In particular, there seem to be two live interpretations of Kampa's theory. According to what I call the *unqualified interpretation* of NS, type admissibility and I-property graining are unqualified in the following sense:

Unqualified Admissibility (UA)

Admissible types can be partially defined by *any* tri-level properties.

Unqualified I-Property Graining (UG)

I-properties can be more or less-fine grained with respect to *any* phenomenal properties of inputs and/or *any* environmental properties of inputs.

On the other hand, according to the *qualified* interpretation of NS, type admissibility and/or I-property graining are qualified in the following ways:

#### Qualified Admissibility (QA)

Admissible types can be partially defined only by tri-level properties of kind X.

#### Qualified I-Property Graining (QG)

I-properties can be more or less-fine grained only with respect to either

- i. input phenomenal properties of kind Y
- or
- ii. input environmental properties of kind Z.

For simplicity, any version of NS that incorporates *either* QA or QG counts as a qualified interpretation of NS.

As I'll argue below, NS faces the following dilemma. The unqualified interpretation would entail that NS countenances relevant types that are far too narrow, thus delivering implausible justification verdicts on particular cases of belief formation. On the other hand, if we adopt the qualified interpretation, then NS fails to make substantive progress towards solving the generality problem by leaving crucial questions unanswered—questions that simply re-raise the generality problem in a slightly different form. As a result, NS either succumbs to counter-example, or it offers at best negligible progress towards solving the generality problem.

### 4.1 The unqualified interpretation and descriptive narrowness

To begin, UA—as opposed to QA—appears to be the most natural way of interpreting Kampa's admissibility condition. After all, in his explicit presentation of admissibility, Kampa places no limitations on the kinds of tri-level properties that can partially define admissible types (236). Secondly, UG—rather than QG—seems to be the the most natural way of interpreting Kampa's account of I-property graining. As I mentioned above, Kampa doesn't offer us a detailed account of I-property graining that covers all the specific ways in which I-property graining can occur.<sup>12</sup> Instead, Kampa gives the examples of phenomenal property graining and environmental property graining to be *illustrative* of how I-property graining works. This being the case, given that Kampa doesn't flag even the possibility that some instance of phenomenal (or environmental) graining might fail to ground an instance of I-property graining, it seems reasonable to interpret NS in terms of UG.

However, if we interpret NS according to UA and UG, then NS produces counterintuitive justification verdicts on particular cases. To see this, let's consider a case from Dutant and Olsson's criticism of Beebe's original tri-level statistical solution to the generality problem.<sup>13</sup> Dutant and Olsson ask us to consider Smith, who uses a very odd cognitive procedure for

---

<sup>12</sup> In one way, this is understandable, as the presentation of such an account likely warrants its own paper.

<sup>13</sup> This case occurs in the section of the paper where Dutant and Olsson are considering various strategies one might use to repair TS to avoid the trivialization problem (1357-1359). This case helped illustrate the repair strategy of "loosening the homogeneity" condition in TS by building in that relevant proper sub-types must be "determined *without* reference to" the truth/falsity outcomes of the tokens comprising that proper sub-type (1358). However, even with this restriction built in to TS, Dutant and Olsson show that TS would still countenance relevant types that are far too narrow, delivering justification verdicts that are intuitively implausible. As I noted above, Kampa himself agreed that this repair strategy fails to salvage TS (235).

identifying tree species throughout his afternoon hike. This procedure can be roughly described as follows: “Smith classifies any tree whose leaves are biggish as a maple tree” (1358). For simplicity, let’s use the tri-level properties  $I_s$ ,  $M_s$ , and  $S_s$  to denote Smith’s odd cognitive procedure used throughout his forest walk. Intuitively, Smith’s tokens that instantiate  $[I_s, M_s, S_s]$  produce beliefs with a very low degree of justification. After all, assuming that Smith is hiking through a normal forest—filled with many non-maple trees that nonetheless have “biggish” leaves—the  $[I_s, M_s, S_s]$  procedure ends up yielding many false beliefs for Smith throughout the afternoon.

However, it does not seem as if NS, once coupled with UA and UG, can accommodate the intuitive verdict on Smith’s case. For simplicity, let  $T_B$  denote a broad type partially defined by  $[I_s, M_s, S_s]$ . Further, let  $t_m$  denote one of Smith’s  $T_B$ -instantiating tokens from that afternoon in which Smith actually happened to be looking at a real maple tree. As it turns out, we can isolate sub-types of  $T_B$  that are both instantiated by  $t_m$  and admissible so long as we assume UA and assume that I-property graining can occur with respect to *any* phenomenal or environmental properties of input experiences. Importantly, these sub-types are also *very* reliable (and hence, statistically relevant), which leads NS to deliver an implausible justification verdict for the belief produced by  $t_m$ . In what follows, I’ll first present the admissible type due to phenomenal I-property graining and then present the admissible type due to environmental I-property graining. After that, I’ll argue that both such types are very reliable according to the most popular conceptions of reliabilism on offer.

First, there is an admissible sub-type of  $T_B$  that emerges due to phenomenal graining so long as we assume UG and UA. Presumably, Smith looks at all kinds of “biggish-sized” leaves during his afternoon hike: maple leaves, birch leaves, beech leaves, oak leaves, etc. Clearly, maple leaves *look* differently than these other leaves. They have a noticeably unique shape. In other words, the phenomenal properties constitutive of the input experiences had while looking at an actual maple leaf are different than the phenomenal properties constitutive of the experiences had while looking at these other kinds of leaves. Let  $L$  denote the set of phenomenal properties that uniquely characterize visual experiences produced when looking at objects that are in fact maple leaf-shaped.<sup>14</sup> Given UG, phenomenal property  $L$  determines an I-property that is finer-grained than  $I_s$ —call it  $I_{L,s}$ . Next, let  $T_L$  denote a type instantiated by  $t_m$  that is partially defined by the cognitive procedure  $[I_{L,s}, M_s, S_s]$ . Given UA,  $T_L$  is also an admissible subtype of  $T_B$ .

Furthermore, assuming UA and UG, there is an admissible sub type of  $T_B$  that emerges due to environmental property graining. Of course, there are many different sorts of descriptions one could give of the environmental conditions under which input experiences

---

<sup>14</sup> Importantly, Kampa adopts Austen Clark’s account of sensory phenomenal properties as “a quality that qualifies the way things appear” (Clark 407). For more detail, see fn. 31 of Kampa (2018: 242). In another footnote, Kampa suggests that the phenomenological properties “associated with I-properties are *post-concept-application*” properties (242, fn. 30). As Kampa explains, a subject  $S$  possesses concepts of phenomenological properties  $A$  and  $B$  just if  $S$  *can distinguish in thought* (by using one’s senses) between something’s exemplifying  $A$  and something’s exemplifying  $B$  (242). For the sake of argument, I’ll assume that this account of phenomenological properties and concept possession is correct. This makes no problem for my example, as it’s reasonable to assume that Smith possesses the concept of  $L$ , i.e., he can tell when something has the  $L$  shape and when something has a shape distinct from  $L$ . Of course, what Smith doesn’t know is that  $L$  happens to correspond to the unique shape of maple leaves.

are produced. That being said, the following description seems to genuinely characterize one such environmental feature that an input experience could have:

- K Having been produced while the subject’s eye lens and retina are focused on an organism whose cells have DNA of type  $k$ —where  $k$  is the sort of DNA that only maples in fact have.<sup>15</sup>

Given UG, environmental property  $K$  determines an I-property that is finer-grained than  $I_s$ —call it  $I_{Ks}$ . Let  $T_K$  denote a type instantiated by  $t_m$  that is partially defined by the cognitive procedure  $[I_{Ks}, M_s, S_s]$ . Given UA,  $T_K$  is also an admissible subtype of  $T_B$ .

As it turns out,  $T_L$  and  $T_K$  are *very* reliable according to the two most common approaches to determining reliability measurements for relevant process types. Let’s call the first such theory *actual-world reliabilism*, according to which the truth-ratio measurement for a given type  $T$  is taken across the class of all  $T$  instances *from the actual world*.<sup>16</sup> Given that  $T_L$  is partially defined by  $I_{Ls}$ ,  $T_L$ ’s extension is only comprised of instances of visual *maple tree* identification where the subject happens to be experiencing the unique phenomenology caused by maple tree leaves. As a result, the truth ratio across all such process instances from the actual world will be quite high. Similarly, the extension of  $T_K$ —given that it’s partially defined by  $I_{Ks}$ —is comprised of instances of visual *maple tree* identification in which the input experience happens to be produced while the subject’s eyes were focused on objects with maple tree DNA. Clearly, the truth ratio across all of the actual world instances of  $T_K$  will be very high (and perhaps perfect).

The other popular approach to measuring reliability is what I’ll call *nearby-worlds reliabilism*.<sup>17</sup> According to nearby-worlds reliabilism, the justification-determining truth ratio for a given process type  $T$  is taken across the class of  $T$ ’s instances throughout all of the possible worlds that are counterfactually close (enough) to the actual world.<sup>18</sup> Given contemporary counterfactual semantics, the notion of being “counterfactually close enough” is something of a vague term of art.<sup>19</sup> That said, while we might not be in a position to offer an exhaustive

---

<sup>15</sup> See Dutant and Olsson (2013: 1357-60). Dutant and Olsson describe how typing Smith’s belief-forming process with respect to this DNA property shows that the “loosening homogeneity” repair strategy for TS (which forbids typing by reference to truth/falsity outcomes) fails to avoid the result that TS countenances relevant types that are too narrow (1359).

<sup>16</sup> Originally, Goldman called this a *frequentist* approach to reliability measurements (1979).

<sup>17</sup> See Goldman (1988: 62-3) for a defense and explanation of this counterfactual approach to reliabilism. In Goldman (1979), he calls this counterfactual approach to measuring reliability a “propensity” account of reliability.

<sup>18</sup> See Stalnaker (1968) and Lewis (1973) for a thorough treatment of possible worlds semantics for counterfactual sentences.

<sup>19</sup> While philosophers haven’t produced anything like an informative analysis of this counterfactual similarity (or, “ordering”) relation, some, like David Lewis, have argued that there are various “weights and priorities” for how various features of a world  $W$  determine how far or close  $W$  is to the actual world  $@$  (1979: 465, 472). Lewis suggests that broad consistency in the laws of nature is the weightiest feature for determining the counterfactual closeness of  $W$  and  $@$ . Where  $L$  is a law of nature in  $@$ , *ceteris paribus*, if  $W$  contains a widespread violation of  $L$ ,  $W$  would be very distant from  $@$ . The second weightiest consideration determining the closeness of  $@$  and  $W$  is the extent (size) of the spatiotemporal regions of  $@$  and  $W$  in which the same facts obtain (472). *Ceteris Paribus*, the smaller the spatiotemporal regions in which  $@$  and  $W$  share the same facts, the further  $W$  is from  $@$ . Also, see Pearl (2000: 239) for a slightly different approach to counterfactual semantics.

description of the counterfactually nearby modal space for any given possible world, few philosophers doubt that we have at least some grasp of this concept—thus making it useful in many contexts.

To begin, it's instructive to consider what a possible world  $W$  would need to be like in order for  $T_L$  or  $T_K$  to be *unreliable* across the nearby possibility space to  $W$ . For  $T_L$ , we can imagine this possible state of affairs:

#### DISGUISE

All of the relevant details are the same as the original Smith-hiking case except for this key alteration: unbeknownst to Smith, a local team of tricksters have a longstanding habit of removing virtually all of the real maple trees in Smith's environment. In addition, the tricksters have painstakingly gone through the entire forest and disguised the poplar trees to look like maples in the following fashion: using scissors, the tricksters have carved all of the poplar leaves into the exact shape of maple leaves.

Let  $W_D$  denote a possible world that satisfies the description of DISGUISE. Notice, in  $W_D$ , some of Smith's token processes instantiate  $T_L$ , given that he'll still experience  $L$  phenomenology when looking at the altered poplar leaves. This being the case, it follows that type  $T_L$  generates mostly false beliefs in  $W_D$ . Given that the worlds counterfactually close to  $W_D$  only include slight alterations from  $W_D$ , it's reasonable to assume that  $T_L$  has a low truth ratio across these nearby possibilities as well.

Let  $@$  denote the actual world in which token  $t_m$  occurs. For our purposes, it's crucial to note that  $W_D$  is, counterfactually, very far away from  $@$ . DISGUISE involves radical divergences from the actual state of affairs in  $t_m$ . Moreover, I contend that cases like DISGUISE are instructive in a more general sense: they highlight just how different the world *would have* needed to be in order for most instances of  $T_L$  and  $T_K$  to have generated false output beliefs. Given the features of  $@$ —and given that  $@$ 's nearby worlds feature only minute differences from  $@$ —it seems that the vast majority of  $T_L$  and  $T_K$  instances that are counterfactually close to  $@$  are ones in which Smith makes *maple*-ascribing judgments about things that are genuine maples.<sup>20</sup> As a result, it's reasonable to conclude that  $T_L$  and  $T_K$  are very reliable according to nearby worlds reliabilism.

As we've seen, arguably the two most influential ways of understanding reliability measurements yield the untoward result that  $T_L$  and  $T_K$  are very reliable and hence statistically relevant partitions of  $T_B$ .<sup>21</sup> Given that these types are also admissible according to UA, it

<sup>20</sup> Of course, this is consistent with there being *some* nearby (to  $@$ ) possible instances of  $T_L$  and  $T_K$  that generate false output beliefs (perhaps due to luck, quantum indeterminacy, or other factors). However, given features of  $@$ , it's reasonable to assume that the false belief-producing instances of  $T_L$  and  $T_K$  constitute a *small measure* of the total nearby possibility space to  $@$ .

<sup>21</sup> There are other approaches to measuring reliability one *could* adopt so as to avoid this result for  $T_K$  and  $T_L$ . Consider what we might call *all-worlds reliabilism*, according to which a type  $T$ 's reliability measurement is taken across every metaphysically possible instance of  $T$ . This being the case,  $T_L$ 's truth ratio across worlds like  $W_D$  would determine whether  $t_m$  generates a justified belief. However, virtually no reliabilists subscribe to all worlds reliabilism, and for good reason. First, it's unclear why the track record of one's belief-forming process in far off possible worlds, including worlds where one is radically deceived by a demon, should matter to the justification of her beliefs formed in our normal non-demon world. Secondly, it's unclear how we would even

follows that they are genuine candidates for being the relevant types for  $t_m$  according to Kampa's NS theory of type relevance. But this is clearly the wrong result, as the belief produced by  $t_m$ , intuitively, has a very low degree of justification.

It looks as if NS succumbs to a very similar problem as TS—determining relevant types that are actually far too narrow. Plausibly, the extension of  $t_m$ 's relevant type should include instances of Smith's odd belief-forming procedure from the afternoon hike in which he incorrectly ascribes *maple* to non-maples. However, NS—according to UA and UG—cannot deliver this result.

Here, one might reasonably think that my narrowness-based counter-example to NS emerges simply because reliabilism *in general* suffers from a much broader objection—namely, the idea that process reliability is *insufficient* for justification.<sup>22</sup> The most famous case marshalled in defense of the insufficiency objection is Laurence Bonjour's *Norman the Clairvoyant* scenario. In this case, Norman finds himself forming beliefs about the far off location of the president of the United States, and as it turns out, these beliefs are always correct (1985: 41). Norman, however, “has no evidence” in favor of these beliefs nor evidence to believe that he has a reliable clairvoyant power (41-42). As Bonjour describes, “[f]rom [Norman's] subjective perspective, it is an accident that his belief [about the president] is true” (43). By Bonjour's lights, this is a straightforward case of reliable, yet unjustified, belief formation (42).

However, there is a key structural difference between my narrowness objection to NS and the insufficiency objection to reliabilism. The narrowness objection notes that reliabilism, as interpreted according to NS, UA, and UG, generates implausible justification verdicts *because* it identifies intuitively *unreliable* cases of belief formation as being highly reliable.<sup>23</sup> On the other hand, the insufficiency objection to reliabilism asserts that there are intuitively clear cases of reliable belief formation that nonetheless result in unjustified belief. Due to this structural difference, there are two response strategies open to the reliabilist with respect to the insufficiency objection that aren't open to Kampa with respect to the narrowness objection.

In response to cases like Norman, reliabilists could, first, attempt to argue against the additional internalist necessary conditions on justification that Bonjour suggests.<sup>24</sup> Secondly, reliabilists could possibly *deny* that Norman's process is a clear case of reliable belief formation. At this point in the dialectic, there is no decisive and widely accepted solution to the generality

---

begin to evaluate the justificatory status of our beliefs, given that we have at best only a faint grasp of the total distribution of *all* possibility space. Similarly, consider *transglobal reliabilism*, according to which the reliability of  $T$  is measured across  $T$ 's instances occurring in environments that are “experientially possible” for the subject, i.e., possible environments where the subject has experiences like the ones the subject *actually* has (Henderson and Horgan 2011). On this view, possible hallucinatory BIV-induced L-experiences and scenarios like DISGUISE would constitute part of the reliability measurement for  $T_L$ . However, transglobal reliabilism is highly controversial. Graham (2014: 531-4) argues forcefully that transglobal reliability is too demanding and thus *unnecessary* for justification.

<sup>22</sup> I'm grateful to an anonymous referee for highlighting these salient connections between the insufficiency objection to the reliabilism and the narrowness objection to NS.

<sup>23</sup> Dutant and Olsson specifically stipulate that Smith's procedure is “very unreliable” as they construct the scenario (2013: 1358).

<sup>24</sup> On this topic, Michael Bergmann has argued at length that adopting such internalist necessary conditions—that are often suggested as a result of considering the case of Norman—lead to vicious epistemic regress problems (2006: 3-24).

problem. Hence, it's open to the reliabilist to suggest that, according to the correct theory of type relevance, it *could* turn out that the relevant type for Norman's process is something like [forming a belief in a way that seems accidental from the subject's perspective]. At the very least, it's by no means clear that this type is reliable.

Notice, neither of these two response strategies is available to Kampa with respect to defending NS from the narrowness objection. This is precisely because Kampa is engaged in offering a particular answer to the generality problem on behalf of reliabilism. I have argued that intuitively unreliable instances of belief formation (like Smith's) turn out to be highly reliable according to NS once interpreted according to UA, UG, and the most common approaches to measuring reliability. Crucially, this needn't be a problem that plagues *any* response to the generality problem, as it's clearly the case that the unique features of Kampa's view—like admissibility and objective homogeneity—are what determine the overly narrow type assignments of NS according to the unqualified interpretation.<sup>25</sup> Responses to the generality problem that omit these factors needn't suffer the same fate.

## 4.2 The qualified interpretation and re-raising the generality problem

As the previous discussion illustrated, the Smith case functions as a counter-example to NS only on the assumption that *both* UA and UG are the correct interpretations of Kampa's view. For instance, if QG is correct, then we cannot just assume that I-property graining can occur with respect to either phenomenal property L or environmental property K. For all that's been said, L might not fall under phenomenal property kind Y and K might not fall under environmental property kind Z. Also, if QA is part of the correct interpretation of NS, then we cannot assume that  $T_L$  or  $T_K$  are admissible sub-types. While I noted that an unqualified interpretation of Kampa's theory appears to be more natural than a qualified interpretation, I needn't commit to a particular interpretation here as I present my dilemma. In what follows, I only argue that *were* QA or QG correct interpretations of NS, then NS fails to make substantive progress towards solving the generality problem.

To see this, consider once again the original puzzle raised by the generality problem. Without some principled way of typing or categorizing a belief-forming process token, it's unclear how to determine whether a given belief has justification according to reliabilism. On

---

<sup>25</sup> A defender of NS—coupled with UA and UG—could use another tactic to accommodate our intuitions on the Smith case and  $t_m$ . She could grant that Smith has first order justification for her belief that  $p$  ( $x$  is a maple), but capture the oddness of this case by denying that Smith has second order justification for believing  $p$ . Upon reflection, I'm inclined to think that this proposal faces a serious challenge in how the counter-intuitiveness of granting justification to Smith at the first order can easily re-emerge at the second order given the details of NS, UA, UG, and how we could fill in Smith's case. From a reliabilist perspective, having second order justification for believing  $x$  is a maple only requires having a reliably formed belief that *my belief that 'x is a maple' is justified*. Similar to how Smith uses an absurd procedure to form his first order beliefs about maples, we could also imagine that Smith uses an equally absurd procedure to form his beliefs about which of his beliefs are justified. But given UA and UG, it could be the case that there's a phenomenally or environmentally finer-grained sub type of this procedure that happens to be highly reliable (and thus, statistically relevant). But this would just be a case where NS delivers counterintuitive justification verdicts for Smith at *both* the first and second orders. Much thanks to an anonymous referee for raising this response strategy.

the assumption that only *types* (rather than tokens) can be measured for reliability, we need some clear conception of a token's relevant type in order to assess the resultant belief's degree of justification. Such a clear conception is what we need from an acceptable theory of type relevance. Reliabilists can't rest content in leaving the notion of *relevant type* undefined or unexplained in their theory.

Now, let's assume that QG is the correct interpretation of NS. In this case, NS would leave two crucial notions undefined in the theory: phenomenal properties of kind Y and environmental properties of kind Z. Now, there's nothing *inherently* problematic with leaving a key notion undefined in a philosophical theory. For instance, the reasonability of a reliabilist theory of justified belief doesn't seem to depend on whether reliabilists can offer a supplementary account of *belief*. However, a theory of justification that invokes kind Y or kind Z without further analyzing these notions seems to raise the exact same sort of problematic ambiguity that's raised by any version of reliabilism that leaves *relevant type* undefined. By leaving these notions unanalyzed, we don't know how to apply the theory in particular cases to see whether it delivers intuitively plausible justification verdicts.

For example, Kampa claims that environmental property graining allows NS to yield the correct justification verdict on cases of visual belief formation under environmental conditions that are "favorable" (2018: 239-40). Moreover, Kampa suggests that *being formed under favorable conditions* can actually constitute part of the relevant type for such token. According to NS, this is because the I-property partially defining this relevant type is determined by a set of inputs that all share a particular sort of environmental condition description. But what exactly constitutes the *relevant* sense of the "environmental conditions" under which a belief is formed?<sup>26</sup> As we saw with the Smith example, we shouldn't accept that property K can constitute (or partially constitute) the relevant environmental conditions for token  $t_m$ . But if K won't do, *which* environmental properties can constitute the relevant environmental conditions for a relevant type? This is just another way of asking what kind Z really amounts to. Without analyzing kind Z any further, we're left with a theory of type relevance that, first, claims that the environmental conditions under which a belief is formed determine whether the belief is justified, and second, offers no guidance for how to conceive of or demarcate these environmental conditions. As a result, we're left with little understanding as to how to test NS (and reliabilism more generally) to see if it can deliver plausible justification verdicts.

The same problem emerges for phenomenal property graining. If L cannot constitute (or partially constitute) the phenomenal features of the I-property that partially defines the relevant type, then *which* phenomenal features can? Without an analysis of kind Y, on QG we're left with a theory which tells us that a particular phenomenal feature description partially

---

<sup>26</sup> Connee and Feldman ask us to consider a reliabilist theory that leaves "relevant type" undefined. They liken such an account to a theory that analyzes the concept *winning race horse* in terms of *suitable horse*, while leaving *suitable horse* undefined. They note, "[i]n the absence of further explanation, this use of 'suitable' has no definite content. On its own, the phrase 'the suitable type of horse' tells us nothing about what makes horses win races... Clearly, a general basis for identifying suitability is required for the claim to say more than just that something or other makes each winning horse win its race" (1998: 4). Similarly, one might reasonably worry that a reliabilist theory which claims that "being formed under favorable conditions is conducive to reliability and justification" while leaving "conditions" undefined is just as uninformative as Connee and Feldman's winning race theory sketched above.

constitutes the relevant type description, but then doesn't say what this particular phenomenal feature is. Once again, it still seems that we're quite far from having a testable informative reliabilist theory. Lastly, even if one grants UG but then accepts QA, we're still left with little insight into the workings of NS. Without an analysis of kind X (for tri-level properties), we don't know which finer-grained tri-level properties can partially define relevant types, and hence, we won't know which statistically relevant sub-types can determine whether a belief is justified. It is in this sense that adopting the qualified version of NS comes with the cost of simply re-raising the generality problem.

## Conclusion

At first glance, it may have seemed that Kampa's new statistical solution to the generality problem was a significant improvement over Beebe's original statistical solution. But as we have seen, NS is problematic in its own right. More specifically, I have argued that NS faces the following dilemma. If we adopt the unqualified interpretation of NS, then NS countenances types that are far too narrow. On the other hand, if we adopt a qualified interpretation of NS, then we currently lack the information necessary for discerning whether NS delivers intuitively plausible justification verdicts on particular cases. Either way, in its current form, NS fails to make substantive progress towards answering the generality problem.

## References

- Adler, Jonathan and Michael Levin. (2002), "Is the Generality Problem too General?" *Philosophy and Phenomenological Research*. 65.1: 87-97.
- Beddor, Bob and Alvin Goldman. (2015), "Reliabilist Epistemology", *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2015/entries/reliabilism/>.
- Beebe, James R. (2004), "The Generality Problem, Statistical Relevance and the Tri-Level Hypothesis." *Noûs*. 38.1: 177-195.
- Bergmann, Michael. (2006), *Justification Without Awareness*. Oxford: OUP.
- Bonjour, Laurence. (1985), *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.
- Clark, Austen. (2008), "Phenomenological Properties: Some Models from Psychology and Philosophy." *Philosophical Issues*. 18: 406–25.

- Comesaña, Juan. (2006), "A Well-Founded Solution to the Generality Problem." *Philosophical Studies*. 129.1: 27-47.
- Conee, Earl and Richard Feldman. (1998), "The Generality Problem for Reliabilism." *Philosophical Studies*. 89.1:1-29.
- (2002), "Typing Problems." *Philosophical and Phenomenological Research*. 65.1:98-105.
- Dawson, M. R. W. (1998), *Understanding Cognitive Science*. Oxford: Blackwell.
- (2013), *Mind, Body, World: Foundations of Cognitive Science*. Edmonton: Athabasca University Press.
- Dutant, Julien and Erik J. Olsson. (2013), "Is there a Statistical Solution to the Generality Problem?" *Erkenntnis*.78: 1347-1365.
- Feldman, Richard. (1985), "Reliability and Justification." *The Monist*. 68.2: 159-174.
- Goldman, Alvin. (1979), "What is Justified Belief?" *Justification and Knowledge*. Ed. by G.S. Pappas. Dordrecht: Reidel. Pgs. 1–25.
- (1988), "Strong and Weak Justification." *Philosophical Perspectives*. 2: 51-69
- Graham, Peter. (2014), "Against Transglobal Reliabilism." *Philosophical Studies*. 169: 525-35.
- Henderson, David and Horgan, T. (2011). *The Epistemological Spectrum*. Oxford: Oxford University Press.
- Kampa, Samuel. (2018), "A New Statistical Solution to the Generality Problem." *Episteme*. 15.2: 228-244.
- Lewis, David. (1973), *Counterfactuals*. Cambridge: Harvard University Press.
- (1979), "Counterfactual Dependence and Time's Arrow." *Noûs*. 13.4: 455-476.
- Marr, David. (1982), *Vision*. San Francisco: W.J. Freeman.
- Matheson, Jonathan D. (2015), "Is there a Well-Founded Solution to the Generality Problem?" *Philosophical Studies*. 172: 459-468.
- Pearl, Judea. (2000), *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Stalnaker, Robert. (1968), "A Theory in Conditionals." *Studies in Logical Theory, American Philosophical Quarterly*. 2, 98-112.