Knowledge, Evidence, and Multiple Process Types

Jeffrey Tolly

Penultimate Draft
Forthcoming in *Synthese*. Please only cite published version. Official pre-published version available here:
https://rdcu.be/bpQEC

DOI:
10.1007/s11229-019-02146-4

The generality problem is one of the most pressing challenges for reliabilism. The problem begins with this question: of all the process types exemplified by a given process token, which types are the *relevant* ones for determining whether the resultant belief counts as knowledge? As philosophers like Earl Conee and Richard Feldman have argued, extant responses to the generality problem have failed, and it looks as if *no* solution is forthcoming. In this paper, I present a new response to the generality problem that illuminates the nature of knowledge-enabling reliability. My response builds upon the insights of Juan Comesaña's well-founded solution to the generality problem, according to which relevant types are *content-evidence* pairs, i.e., descriptions of both the target belief's *content* and the *evidence* on which the belief was based. While most responses to the generality problem, including Comesaña's, only posit one relevant type for any given process token, I argue that knowledge-enabling reliability requires a process token to be reliable with respect to *multiple* content-evidence pairs, each with varying degrees of descriptive specificity. I call this solution *Multi-type Evidential Reliabilism* (MTE). After offering a clear formulation of MTE, I conclude by arguing that MTE is sufficiently informative to rebut Conee and Feldman's generality problem objection to a reliability condition on knowledge.

Generality Problem, Reliabilism, Knowledge, Externalism, Evidentialism

[Word Count: 9,835]

1. Introduction

The generality problem is one of the most important challenges for reliabilism.[1] According to reliabilism, a belief b counts as knowledge only if b is formed by a reliable process. With this in mind, the generality problem begins with the following puzzle. Processes can be

[1] See Beddor and Goldman (2015). In their overview of reliabilist epistemology, they see the generality problem to be in the top six "problems," or, "objections" to reliabilism.

thought of as general/repeatable process *types* or as precise, one-off process *tokens*. Plausibly, only the former can be evaluated for reliability or unreliability.[2] As both defenders and critics of reliabilism have noticed, a given process token exemplifies countless types. Consider an example of a normal process token in which someone uses perception to form the belief that *there's a jar on the table*. This process token exemplifies innumerable type descriptions, including [perception], [vision], [vision in good lighting conditions], etc. At this point, one might wonder which of these types determine whether a belief counts as knowledge. In other words, which of these type descriptions are epistemically *relevant*?[3] This question seems particularly pressing for the reliabilist, seeing as how these types all have different degrees of reliability. Without an answer to this question, it's unclear what epistemic implications reliabilism has for particular cases of belief formation.

In this paper, I attempt to make substantive progress towards solving the generality problem. Here, I present a new theory of process type-relevance that illuminates the nature of knowledge-enabling reliability, i.e., the reliability necessary condition on knowledge. Importantly, the generality problem has traditionally been framed as a problem for both reliabilist accounts of knowledge *and* reliabilist accounts of justification.[4] That said, in order to properly limit the scope of this paper, I will remain neutral on the degree to which my theory of type-relevance—which explicates knowledge-enabling reliability—also serves to illuminate justification-enabling reliability.

I call my theory of type-relevance *Multi-Type Evidential Reliabilism* (MTE). MTE builds upon the insights of Juan Comesaña's well-founded solution to the generality problem, according to which relevant types are content-evidence pairs, i.e., descriptions of both the content believed in the token case and the evidence on which that belief was based. However, contrary to Comesaña's well-founded solution, MTE holds that each token has *multiple* relevant content-evidence pair descriptions.

In §2, I present Conee and Feldman's canonical description of the generality problem for reliabilism and how its currently un-answered state constitutes a powerful objection to reliabilism. Over the past three decades, Conee and Feldman (and others) have shown how virtually every response to the generality problem fails to meet what I call *the specificity challenge:* the challenge of typing processes in a way that captures a token's epistemically relevant degrees of descriptive specificity.

In §3, I present Comesaña's well-founded solution. After discussing its relative merits, I show how it fails to meet the specificity challenge in how it only selects descriptively narrow process types as the *sole* relevant types for any given process token. That said, the cases used to demonstrate this result also suggest a reasonable repair strategy: a correct answer to the generality problem must state that tokens generate knowledge only if they're

---

[2] Richard Feldman (1985) and Goldman (1979) are key figures who highlighted this important distinction for making sense of reliabilism. Most philosophers agree that types, rather than tokens, are the entities that can be measured for reliability. Although, Juan Comesaña (2006) articulates a way in which tokens could be evaluated for reliability by taking a truth-ratio for that token across some class of possible worlds. However, Comesaña recognizes that framing reliabilism like this doesn't get the reliabilist out of the generality problem (28-30). The reliabilist still would have to provide an account of *which* possible worlds were included in this truth-ratio.

[3] The specific language of "relevant" and "irrelevant" types was first introduced by Feldman (1985: 160).

[4] For instance, in Conee and Feldman's canonical treatment of the generality problem, they describe the generality problem as undermining both "process reliability theories of justification and knowledge" (1998:24). Interestingly enough, the growing literature on the generality problem does not feature any substantive discussion on the degree to which type-relevance for knowledge-enabling reliability correctly characterizes type-relevance for justification-enabling reliability, or vice versa.

reliable with respect to *multiple* content-evidence pairs, each with varying degrees of descriptive specificity.

In §4, I argue that the counterfactual nature of knowledge-determining reliability measurements provides a fitting model for describing a token's numerous relevant content-evidence pairs. There, I formulate MTE by invoking notions like *similarity relations* and *degrees of difference/dis-similarity*. In §5, I conclude by arguing that MTE is sufficiently informative so as to successfully undermine Conee and Feldman's generality problem objection to a reliability condition on knowledge.

2. The Generality Problem Objection

2.1 The Specificity Challenge

Famously, Earl Conee and Richard Feldman claim that, "Without a specification of the relevant type, process reliabilism is radically incomplete. Only when a bearer of reliability has been identified does the theory have any [epistemic] implications about…beliefs in particular cases" (1998: 3). They continue, "[Reliabilism] must first be elaborated at least enough to imply *exactly* what process type has to be reliable in the case in question" (3). According to Conee and Feldman, reliabilists have a burden to elaborate their theory by giving a supplementary account of type-relevance that is principled: it cannot "make *ad hoc* case-by-case selections of types that match our intuitions" (4). In addition to being principled, it must also be "universal" in the sense that it must "specify the relevant type whenever there are definite facts" about whether a given process generates knowledge (4).

In addition to offering an account of type-relevance, Conee and Feldman also claim that reliabilists must "specify which situations of a process type's operation determine whether or not the type is reliable" (3). This relevant class of situations is what statisticians call the *reference class* for truth-ratio measurements.[5] For example, suppose that a given token t has a relevant type T of [visually-based color judgment].[6] We can sensibly ask, *"Across which* class of visually-based color judgments is the knowledge-determining truth-ratio measured?" All of the visually-based color judgments of the subject's lifetime? All of the visually-based color judgments throughout her society's history? All of the counterfactually nearby visually-based color judgments? Of course, there are other options, but clearly this question is of central importance given that each of the classes mentioned above will (or can) have different truth-ratios.[7] So, by Conee and Feldman's lights, any acceptable solution to the generality problem must be principled, universal, and supply an informative account of both type-relevance and reference classes. For the purposes of this paper, I accept this description of the generality problem and the corresponding argumentative burden that it places on the reliabilist.

In addition to *stating* the generality problem for reliabilists to solve, Conee and Feldman go one step further. Famously, they argue that the currently un-solved state of the generality

[5] See Reichenbach (1949: 374) for a presentation of reference classes and their role in probability theory.
[6] Throughout the paper, I use the term "judgment" interchangeably with "belief."
[7] See Wallis (1994) for a response to the generality problem that features a helpful discussion of the process type/reference class distinction. Note, Wallis calls reference classes "relevance classes" in his article.

problem constitutes a compelling reason to reject reliabilism. I call this argument of theirs *the generality problem objection* to reliabilism.

The generality problem objection begins with the following observation. Conee and Feldman—as well as other philosophers—have shown there to be serious problems with nearly every extant response to the generality problem.[8] Given the failure of these various proposals, Conee and Feldman assert that "[t]here is no significant progress in any of these approaches, singly or in combination. The basic process reliabilist idea just does not pan out" and that "the prospects for a solution to the generality problem for process reliabilism are worse than bleak" (1998: 5). Furthermore, after criticizing many of these responses to the generality problem, Conee and Feldman draw the following gloomy conclusion:

> That is the full variety of existing approaches to disposing of the generality problem. In the absence of a brand new idea about relevant types, the problem looks insoluble. Consequently, process reliability theories of justification and knowledge *look hopeless*. (24, emphasis mine)

Importantly, Conee and Feldman do not offer a detailed explanation for *why* this failure to solve the generality problem shows that reliability theories of knowledge are hopeless. Offering a thorough treatment of how to best reconstruct their reasoning is a valuable project in its own right—a project that is unfortunately outside the scope of this essay. That said, in what follows I'll briefly present what I take to be a plausible formulation of Conee and Feldman's argument, and this should suffice for our purposes here.

To begin, from these quotations we can sensibly interpret Conee and Feldman as suggesting that the long track-record of failure to solve the generality problem constitutes strong inductive evidence that there *probably isn't any solution* to be found. This result would pose a serious problem for a reliability theory of knowledge. As I discussed above, a reliability theory of knowledge posits a reliable-process necessary condition on knowledge. Whether a given belief satisfies the reliability condition crucially depends on whether the relevant types corresponding to the belief-forming process token are sufficiently reliable. But, if there probably isn't any correct theory of type-relevance to be found, then there probably aren't any facts about which types are relevant for any given token.[9] If there probably aren't any facts

---

[8] See Conee and Feldman (1998) for a criticism of most of the relevance theories presented up through that point in time, with a particular focus on Heller (1995), Alston (1995), Schmitt (1992), and Goldman (1979). For other theories of type-relevance, see Beebe (2004), Leplin (2007), Wallis (1994), Becker (2008) Adler and Levin (2002), Lepock (2009), Sosa (1991), Goldman (1986), Comesaña (2006), Greco (2010), Olsson (2016), and Wallbridge (2016). Also see Conee and Feldman (2002) for an updated criticism of Adler and Levin's (2002) relevance theory. See Brueckner and Buford (2013) for a criticism of Becker (2008), and Dutant and Olsson (2013) for a criticism of Beebe (2004). See Matheson (2015) for a criticism of Comesaña (2006).

[9] Conee and Feldman's demands "to know what *determines* [type] relevance" and insistence that this determining factor is "principled" rather than "ad-hoc" indicate that they are assuming the following claim: *if* there are facts about type relevance, then these facts must be explained or determined by some set of general principles (1998: 3-4, emphasis mine). While this is the dominant assumption throughout the generality problem literature, not everyone agrees that type-relevance must have some *principled* and general theoretical account or explanation. Klemens Kappel defends what he calls the "no determination view," according to which "there are no facts that determine relevant types" for particular tokens (2006: 256). In other words, the characteristics of a token's relevant type are *unexplained* by any broader theoretical principles. It's beyond the scope of this paper to offer a full rebuttal to Kappel. In what follows, I'll agree with the majority of contributors to the generality problem literature and assume that *if* there are facts about which types are relevant for any given token, then there exists some is a reliability condition on knowledge, then type-relevance has a principled account or analysis.

about which types are relevant, then there probably aren't any facts about whether particular beliefs satisfy the reliability condition on knowledge.[10] However, on the assumption that there are (at least some) facts of the matter about which beliefs count as knowledge, we should then infer that there probably isn't a reliable-process necessary condition on knowledge.

Here, we should focus on the first—and most substantive—step in this argument: why might the history of failure to formulate an account of type relevance constitute good reason to believe that there's most likely *no such account* to be found? To answer this question, it's important to grasp the particular way in which most extant generality problem responses fail. First off, Conee and Feldman are quick to point out that some extant responses are insufficiently filled out and as a result lack any entailments about which particular types are relevant in various cases. Thus, these responses fail the "universality" requirement on solving the generality problem.[11] But for the responses that successfully generate type-relevance verdicts for cases of belief-formation, Conee and Feldman adeptly show how these responses either type processes too narrowly or too broadly. Consequently, these responses entail implications for knowledge-enabling reliability that are intuitively implausible.[12]

In his earliest paper on the generality problem, Richard Feldman predicted that the following trap awaited any attempt to solve the generality problem: to the extent that a given theory of type-relevance avoids typing processes too narrowly, it will type them too broadly; to the extent that a theory of type-relevance avoids typing processes too broadly, it will type them too narrowly (1985: 160-1). Given the subsequent decades of failure to solve the generality problem, it would appear that Feldman's early skepticism has been vindicated.

That said, we can more precisely describe the failure of previous generality problem responses as a failure to meet the following challenge:

The Specificity Challenge:	Formulate a principled account that correctly captures the relevant degrees of descriptive type specificity so as to deliver intuitively correct epistemic verdicts on particular cases of belief formation.

[10] The way that I reconstruct Conee and Feldman's reasoning here represents the inductive evidence (of failure to solve the generality problem) as supporting a *metaphysical* conclusion: there probably *are no facts* of the matter about type relevance and knowledge-enabling reliability. But one could interpret the inductive evidence to have more of an epistemic upshot: *we could never know or reasonably judge* any facts about type relevance and knowledge-enabling reliability. This interpretation of the argument would still (if cogent) undermine our reasons for accepting a reliability condition on knowledge. After all, if we could never know or reasonably determine any facts about knowledge-enabling reliability, then one should worry about our ability to apply reliabilism to particular (actual or hypothetical) cases of belief formation in order to test reliabilism for explanatory power—i.e., to test whether reliabilism accommodates our epistemic intuitions on a variety of cases.

[11] For example, Conee and Feldman (1998: 6-11) argue that the "common-sense types" and "natural kinds" responses to the generality problem both fail to deliver a verdict on which types are relevant for any given token.

[12] For instance, Conee and Feldman show that Alston's theory types processes too narrowly (1998: 13-15) and that Adler and Levin's theory (2002) types processes too broadly (Conee and Feldman 2002: 102-3). Moreover, Brueckner and Buford (2013) show that Becker's (2008) theory types processes too narrowly, and Dutant and Olsson (2013) show that Beebe's (2004) theory types processes too narrowly. To clarify, narrower type descriptions build in more detail from the token process, and broader type descriptions build in less detail from the token process.

According to Conee and Feldman, the generality problem objection gains its plausibility because, after considering a wide variety of responses, it seems unlikely that *any* principled proposal *could* escape Feldman's trap and meet the specificity challenge.[13] Furthermore, if it's unlikely that there's any working theory of type relevance to be found, then it's unlikely that there's any reliability condition on knowledge.

2.2 Framing the Generality Problem

Given the track record of failure to formulate a correct theory of type-relevance, the generality problem objection appears to be quite formidable. Nevertheless, I think there's good reason to believe that the *apparent* intractability of the specificity challenge stems from a subtle flaw in the way that the generality problem is commonly presented. Consider how Conee and Feldman initially frame the burden that reliabilists have for solving the generality problem: "[W]ithout a specification of *the* relevant type, process reliabilism is radically incomplete. Only when *a* bearer of reliability has been identified does the theory have any implications about the justification of beliefs in particular cases" (1998: 3 emphasis mine). Furthermore, consider the grounds on which they criticize one of the extant theories of type-relevance: the natural kinds theory. The natural kinds theory states that the "relevant type for any belief-forming process token is the natural kind to which it belongs" (10). Conee and Feldman criticize this theory by noting that, "Process tokens may belong to natural kinds. Still, there is no good reason to think that each token belongs to just *a single* natural kind, and hence no reason to think that [the natural kinds theory] provides a solution to the generality problem" (10, emphasis mine).

These passages seem to suggest that Conee and Feldman are presupposing that knowledge is determined by the reliability of *a single* relevant type description.[14] But this is a substantive assumption about the workings of knowledge-enabling reliability. Why couldn't knowledge-enabling reliability depend on truth-ratios measured across multiple types? Conee and Feldman provide no argument for their single-type rendering of epistemic reliability. Not surprisingly, most of the philosophers who have responded to the generality problem—as laid out by Conee and Feldman—have operated under the assumption that knowledge (and justification) depend on the reliability of just a *single* type.[15]

---

[13] Erik Olsson (2016) offers a similar explanation for why the generality problem objection is so pressing. Given Conee and Feldman's arguments, it seems as if responses to the generality problem will make process types either "too broad" or "too narrow" for delivering justification verdicts that match our "everyday concept of justification" (181).

[14] Also, a single-type approach to epistemic reliability appears to be framing Richard Feldman's original (1985) presentation of the generality problem:

> Let us say, then, that for each belief-forming process token *there is some "relevant" type* such that it is the reliability of *that type* which deter mines the justifiability of the belief produced by that token. Thus, the reliability theory can be formulated as follows:
>
> > (RT) S's belief that is justified if and only if the process leading to S's belief that is a process token whose relevant process *type* is reliable. (1985: 160, emphasis mine)

[15] That said, some multi-type responses to the generality problem have been proposed. The earliest multi-type approach comes from Mark Wunderlich, as he presented his theory of *vector reliability* (2003). Wunderlich claims that *every type* exemplified by the token process plays some role in determining the degree of justification that a resultant belief has (243-5). While Wunderlich explores various approaches for aggregating all of the truth ratios from all of a token's types in order to determine whether that token generates knowledge, he

However, framing the generality problem as a challenge to identify a *single* type for each token profoundly mischaracterizes the nature of knowledge-enabling reliability. In what follows, I present several cases that together support the following conclusion: in order for a token to generate knowledge, it must be reliable with respect to *multiple* types—each with differing degrees of descriptive specificity.[16] As a result, the only way to meet the specificity challenge is to adopt a multi-type approach to solving the generality problem.

Despite being a single-type theory as stated, Juan Comesaña's *well-founded* solution to the generality problem provides us with a plausible framework for conceiving of a token's multiple relevant types. In the next section, I'll briefly present the well-founded solution before turning to my argument for the multi-type nature of knowledge-enabling reliability.

3. The Well-Founded Solution and Multiple Process Types

3.1 Comesaña's Well-Founded Solution to the Generality Problem

Comesaña's well-founded solution takes its cue from William Alston's contention that the *psychological* features of a subject—leading up to belief-formation—crucially determine whether a belief counts as knowledge (Alston 1995: 17). Comesaña interprets Alston as defending the following idea:

> [T]here is a fact of the matter about what psychological features in the mental life of the subject had an impact in the formation of the belief in question, and what kind of impact they had. There will always be, then, a process-type-schema, having been produced by such and such mental antecedents in such and such a way, that we can fill out in each case in order to get a specific process-type to assess for reliability (Comesaña 2006: 34)

Comesaña identifies these relevant mental features as the *evidence* on which a subject bases her belief in the token process:

> [W]e have good reasons to believe that any adequate epistemological theory needs to appeal, either implicitly or explicitly, to the notion of a belief's being based on certain evidence…I will argue that that notion is all we need to solve the generality problem. (2006: 33)

---

ultimately refrains from defending any particular aggregation approach, claiming that such a project is "beyond the scope of this paper" (249). Next, Kevin Wallbridge (2016) defends a multi-type approach to knowledge-enabling reliability. Like my own theory of type-relevance, it too draws inspiration from Comesaña's well-founded solution to the generality problem. See fn. 24 for further discussion on Wallbridge's proposal.

[16] In my (forthcoming), I briefly entertain such a multi-type approach to knowledge-enabling reliability. I argued that *if* such a multi-type approach to reliability were correct, then we'd have good reason to believe that a proper function condition on knowledge is explanatorily dispensable. However, there I offered no substantive arguments in defense of this multi-type answer to the generality problem nor any thorough discussion of how to formulate this reliability condition on knowledge. Here, I take up both of these tasks in what follows.

Comesaña sees the evidential-basing relation as the key epistemically-relevant feature in any belief-forming process token.

> Given that there will always be some evidence that the belief is based on, the process that generates the belief will always instantiate a case of the type-schema *producing a belief that p based on evidence E*. (2006: 37)

So, according to the well-founded solution to the generality problem, a token's sole relevant type is its instance of the schema, *producing a belief that p based on evidence E*.[17] In other words, Comesaña thinks that relevant type descriptions are a certain sort of *content-evidence pair* description—a description of the content adopted for belief in the token case combined with a description of the evidence on which the agent bases her belief in the token case. In particular, Comesaña indicates that a token's relevant type is its *maximally specific* content-evidence pair—the specific proposition p believed in the token case paired with the maximally precise description of the evidence E on which the target belief was based. Given Alston's emphasis on *mental* antecedents to belief-formation, we can charitably interpret Comesaña as invoking a mentalist account of evidence: evidence consists of mental states. These might include other beliefs of an agent, as well as experiences like perceptual experience, intuitive seemings, introspective experiences, etc.[18]

Comesaña's presentation of the well-founded solution doesn't include an account of reference classes for knowledge-enabling reliability measurements. However, in Comesaña's earlier work on the nature of knowledge, he defends a *counterfactual* account of knowledge-enabling reliability.[19] In other words, he argues that a token's knowledge determining truth-ratios are taken across classes of possible belief-forming events that are counterfactually close to the (actual) token process event. We can formulate Comesaña's understanding of reference classes as follows:

> RC Where T is a relevant type description for a token process t carried out by subject S, the belief produced by t counts as knowledge only if there is a sufficiently high truth-ratio across a reference class comprised of all possible belief-forming events that are counter-factually close to t in which S undergoes belief-forming processes that exemplify T.

RC nicely captures the plausible idea that knowledge-enabling reliability depends on the subject's belief-forming *dispositions*. Presumably, the relevant sense of "disposition" here corresponds to the sorts of things that the subject *would* do in a certain class of possibilities that are sufficiently similar to the actual case—where these similar possibilities *could have been actualized* with only minimal changes to the actual course of events. Given that counterfactual closeness is a sort of similarity measurement, we can straightforwardly describe this relevant class of similar possibilities as the ones that are *counterfactually nearby* to the token process itself. A counterfactual approach to knowledge-determining reliability measurements is by no means

---

[17] See Comesaña (2010: 584-93) for further explanation and defense of this solution to the generality problem.
[18] Alston explicitly invokes these kinds of experiences as the grounds or "inputs" to relevant belief-forming processes (1995: 17-18).
[19] Comesaña argued for a "diagonal reliability" requirement on knowledge (2002: 261-2). As he makes clear, diagonal reliability is determined by truth-ratios that are counterfactual in nature (258).

a unique proposal, as others have suggested that knowledge crucially involves some sort of counterfactual connection to the truth.[20]

At first glance, the well-founded solution to the generality problem has much to admire. By focusing in on evidential basing, it identifies the sorts of token features that seem relevant to epistemic evaluation. More importantly, the well-founded solution highlights the epistemic importance of *the entire body of evidence* on which beliefs are based. In everyday life, many of our beliefs are based on bodies of evidence that are rather vast and diverse. For example, our beliefs pertaining to the trustworthiness of a given individual's testimony in a particular situation are typically based on a variety of perceptual experiences as well as a vast amount of background information (stored in memory) about that person and context. Plausibly, *all* of this evidence factors in to whether the corresponding belief-forming process is reliable enough to generate knowledge. The well-founded solution straightforwardly incorporates this insight.

3.2 Multiple Relevant Process Types

It bears pointing out that Comesaña explicitly deploys the well-founded solution to the generality problem in order to formulate an account of *justified* belief. In brief, Comesaña holds that a token process generates a justified belief *if and only if* its maximally precise instance of the type schema [believing p, on the basis of evidence E] is reliable (38). Recently, Comesaña's theory of justification has come under attack. For example, Jonathan Matheson has argued that Comesaña's theory fails to supply a sufficient condition on justification.[21] In particular, Matheson shows that Comesaña's approach to justification delivers implausible justification verdicts on cases where the subject comes to believe a "stable proposition" i.e., propositions whose truth values are stable across nearby possible worlds (2015: 465-7).[22]

While investigating the nature of epistemic justification is *not* the topic of this paper, I do think that Matheson's objection to the sufficiency Comesaña's theory of justification helps to illuminate the reliability necessary condition on knowledge. To begin, I accept that Comesaña's well-founded solution to the generality problem highlights *one* necessary condition on knowledge that pertains to reliability: In order for a token to generate knowledge, it must be reliable with respect to its precise content-evidence pair.[23] That said, as I'll argue below,

---

[20] For example, Alvin Goldman (1988:63) argues that epistemic reliability measurements are taken across possibilities that that are sufficiently close to the actual case. Interestingly enough, Duncan Pritchard's account of the safety condition on knowledge is rather similar to RC:

> If a believer knows that p, then in nearly all, if not all, nearby possible worlds in which the believer forms the belief that p in the same way as she does in the actual world, that belief is true (2005: 163).

It appears as if, according to Pritchard, a belief is safe just if there's a sufficiently high truth-ratio taken across a reference class of possibilities—that are sufficiently nearby—in which the subject forms a judgement "in the same way" that she does in the token case. See fn. 27 below for a further discussion on the relationship between safety and knowledge-enabling reliability.

[21] See Wallbridge (2016: 347) for a different sort of counterexample to the sufficiency of Comesaña's theory of justification. Wallbridge uses an altered version of the famous fake-barn case to construct this counter-example. See fn. 24 for further discussion of Wallbridge's theory of type-relevance.

[22] In Matheson's presentation of this problem, he uses stable propositions that are contingently (yet eternally) true, like *the gravitational force is proportional to the inverse-square of the distances between two objects* (2015: 465-7). While the use of contingent stable propositions suffices for highlighting this narrowness worry for the well-founded solution, I think that judgments on necessary truths (which manifest actual and counterfactual stability) illustrate this point even more clearly.

[23] Much thanks to an anonymous referee who insightfully encouraged me to frame Comesaña's well-founded solution as *one* of the reliability necessary conditions on knowledge.

Matheson's objection provides us good reason to accept *additional* necessary conditions on knowledge that pertain to process reliability. In particular, Matheson's objection suggests that tokens must also be reliable with respect to many of its less-specific content-evidence pairs in order to generate knowledge.

First, consider the following case of believing a stable proposition that is a *necessary* truth:

> MATH 1
>
> Peter enjoys practicing simple arithmetic in his head for hours at a time. He'll think about simple formulas (e.g., 2+6, 40 – 7, 20+14, etc.,) and attempt to intuit their values. Sadly, unbeknownst to Peter, three months ago he took a blow to the head that damaged his arithmetical intuition abilities. Now, he has a stable disposition to intuitively judge *incorrect* values for these formulas roughly 90% of the time. That afternoon, roughly 90% of Peter's arithmetical judgments end up being false. Let t_{m1} represent one of these token processes in which he considers the formula 37+14, and just happens to have an intuitive seeming that correctly represents 37+14 as being equal to 51. Then, on the basis of this seeming, Peter comes to believe *37+14=51*.

According to Comesaña's well-founded solution to the generality problem, t_{m1}'s sole relevant type is the following precise content-evidence pair:

> M_p [judging that *37+14=51*, on the basis of an intuitive seeming that represents 37+14 as being equal to 51]

Matheson directs our attention to the fact that the well-founded solution to the generality problem *builds the specific content of the target belief* into the relevant type description (2015: 465-6). In this example, the content of the resultant belief is a necessary truth, i.e., a belief that is true in *every* possible situation. Hence, any reference class of M_p belief-forming events will (necessarily) have a 100% reliability score. Hence according to the well-founded solution, M_p is maximally reliable.

Matheson is keen to point out that tokens like t_{m1} fail to generate justification, despite satisfying Comesaña's necessary *and sufficient* conditions on justification. However, for our purposes here, I contend that t_{m1} fails to generate knowledge. Intuitively, it does not seem that Peter knows 37+14=51 in this case. Furthermore, there is a plausible explanation for Peter's lack of knowledge that readily presents itself: given Peter's cognitive dispositions, t_{m1} falls far short of having knowledge-enabling reliability. More precisely, upon a straightforward reading of the case, t_{m1}'s knowledge-enabling reliability seems to be undermined due to its astounding counterfactual *un*reliability with respect to the following *broader* content-evidence pair:

> M_b [judgment on the value of a given simple arithmetical formula, on the basis of an intuitive seeming representing a value for that formula]

Secondly, consider a case of believing a contingent truth that is nonetheless a stable proposition:

COLOR 1

A long time ago, Susie trained as a paint color-shade identification specialist. She became very adept at seeing and recognizing very fine-grained color shades and distinguishing them from other color shades. Today, she's practicing her color identification skills on the paint chip board at the local paint shop. Sadly, unbeknownst to her, a few months ago, after taking a hit to the head, her visual faculty was damaged and now has the following dispositional profile: 90% of the time, she visually represents objects as being about 10 color shades off from what they actually are. For example, she's disposed to see objects that are actually green 11 as being green 20 or green 21. Hence, when she goes to practice that afternoon, roughly 90% of her color shade identifications are incorrect. Let t_{c1} represent one of these token processes. In t_{c1}, Susie looks at paint chip o_1—that is in fact orange 47—and then just happens to have a visual experience that represents o_1 as being orange 47. On that basis, she then comes to believe *o_1 is orange 47*. Also, unbeknownst to Susie, a mad scientist secretly made paint chip o_1 with his patented "color consistency" chemical—a molecule that repels substances that would otherwise change the color of o_1. As a result, it is physically impossible to change o_1's color without utterly destroying it.

According to the well-founded solution, t_{c1}'s sole relevant type is the following precise content-evidence pair:

C_p [judging that *object o_1 is orange 47*, on the basis of a visual experience that represents o_1 to be orange 47]

C_p builds in to the relevant type description that Susie succeeds in having a visual experience that correctly represents o_1 to be orange 47. Moreover, given that o_1 is made with the color consistency chemical, o_1 is in fact orange 47 in all (or nearly all) nearby possible worlds. It follows that t_{c1}'s reliability score is very high according to the well-founded solution.

But it seems that token t_{c1} falls short of having knowledge-enabling reliability. Furthermore, it seems like t_{c1}'s counterfactual unreliability with respect to the following broader content-evidence pair is a key factor that undermines Susie's knowledge in this case:

C_b [judging an object's particular color shade, on the basis of visual experiences that represent a particular color shade for that object][24]

[24] Importantly, COLOR 1 highlights a way to construct a counter-example to the sufficiency of Kevin Wallbridge's (2016) account of knowledge-enabling reliability. Wallbridge uses a slightly-altered version of the famous *fake barn* case to point out that, at least in some cases, certain portions of the target belief's propositional content are based on only certain portions of the total evidence used for basing the belief. Furthermore, these different portions of evidence work together in a "combinatorial structure" to support the entire target proposition (2016: 349). For instance, consider a token t_r, in which one visually comes to believe that *object x_1 is a red 14 barn*. Token t_r has the following relevant type components. Component C1: [forming a judgment on whether x_1 is red 14, on the basis of a visual experience representing x_1 to be red 14]; Component C2: [forming a judgment on whether x_1 is a barn, on the basis of a visual experience representing x_1 to be barn-shaped] According to Wallbridge, C1, C2, and the combination type C1&C2 must all three be reliable in order for t_r to generate knowledge (349-50). That said, imagine someone who, counterfactually, is very good at recognizing things to be red 14, but terrible at (accurately) visually representing every other color shade. Such a person's t_r token could be counterfactually reliable with respect to C1, C2, and C1&C2, but since it is also

Importantly, the counterfactual unreliability with respect to broader types like M_b and C_b *do* seem to be relevant for whether a token generates knowledge. Assuming that a token must also be reliable with respect to its precise content-evidence pair in order to generate knowledge, these two cases suggest that we should augment the well-founded solution to take a *multi*-type form in order to fully explicate knowledge-enabling reliability.

At this point, it's worth considering *how many* relevant types a given token has with respect to whether the token has knowledge-enabling reliability. As it turns out, there's good reason believe that tokens must be reliable with respect to *numerous* content-evidence pairs in order to generate knowledge. First, consider the following variations of COLOR 1 and MATH 1.

COLOR 2

All details are the same as COLOR 1, except in this case, Susie's blow to the head (unbeknownst to her) gives her the following dispositional profile: when looking at something that is orange 47, her visual faculty will, *most* of the time, represent the object to be a shade that's 10 shades off from orange 47. In addition, when her faculty does represent some object as being orange 47, most of the time the representation is off by 10 shades. However, her cognitive ability to correctly represent all of the other color shades is left untouched and continues to function normally. One day, in the paint shop, she undergoes token process t_{c2}: she looks at o_1 and on this particular occasion, her visual faculty just happens to correctly generate a visual experience that represents o_1 as being orange 47, and on this basis, she comes to believe *object o_1 is orange 47*.

MATH 2

All the details are the same as MATH 1, except in this case, Peter's blow to the head (unbeknownst to him) gives him the following dispositional profile: when considering 37+14, his intuitive faculty will, *most* of the time, generate an intuitive seeming that represents the incorrect value/sum for this formula. However, the blow to the head left his ability to correctly intuit all other simple arithmetical formulas untouched. One day, as he's practicing arithmetic, he undergoes the following token process t_{m2}: he considers 37+14, and on this particular occasion, he just happens to intuit that the sum is 51.

Once again, tokens t_{c2} and t_{m2} don't seem to generate knowledge. In these altered cases, t_{c2} is counterfactually reliable with respect to C_p and C_b, and t_{m2} is counterfactually reliable with respect to both M_p and M_b. However, these tokens are unreliable with respect to the following *narrower* content-evidence pairs (that are nonetheless slightly broader than C_p and M_p):

C_n [forming a judgment on whether a given object is orange 47, on the basis of a visual experience that represents a given object to be orange 47]

M_n [forming a judgment on the value of 37+14, on the basis of an intuitive seeming that represents a value for the sum of 37+14]

counterfactually *un*reliable with respect to C_b (much like Susie's token t_{c1}), it seems like t_r would fall short of having knowledge-enabling reliability.

Furthermore, it seems like t_{c2} and t_{m2}'s unreliability with respect to these narrower content-evidence pairs is sufficient for undermining their knowledge-enabling reliability. Reflecting on these four cases suggests that knowledge-enabling reliability, for any given token, depends on having sufficiently high truth-ratios for *at least three* relevant-type descriptions, each with different degrees of descriptive specificity.

That being said, there appear to be far more than just *three* relevant types for any given token. Consider the following alterations to the previous scenarios:

> COLOR 3
>
> All the details are the same as COLOR 1, except this time, Susie's blow to the head only alters her dispositional profile in this way: her visual faculty will, most of the time, misrepresent (by 10 shades off) every orange color shade *except* orange 47. Her cognitive dispositions to represent all other color shades have been left untouched by the accident. That afternoon in the paint shop, she undergoes a belief-forming process token t_{c3} in which she happens to look at an orange 47 paint chip (o_1), visually represents it as being orange 47, and comes to believe *o_1 is orange 47*. While in this case she *in fact* happened to look at an orange 47 paint chip, leading up to that moment it was much more likely that she would have looked at other shades of orange, given that the paint store has several paint chips for other shades of orange on display as well. Also, the orange-shaded paint chips comprise only a very small percentage of the entire paint shop's color shade display. As a result, the vast majority of Susie's visually-based color shade judgements made throughout the entire afternoon are identifications of non-orange color shades.
>
> MATH 3
>
> All the details are the same as MATH 1, except this time, the blow Peter takes to the head only alters his cognitive dispositions in the following way: were he to consider any simple two-digit addition formula *except for 37+14,* he would most likely have an intuitive seeming that misrepresents the value of that formula. His cognitive dispositions for intuiting the values of every other simple arithmetical formula have remained untouched by the accident. That afternoon, Peter undergoes a token belief-forming process t_{m3} in which he just happens to consider 37+14 and comes to intuit the value as 51. Given that 37+14 is just one among many simple two-digit addition formulas, in the moments leading up to t_{m3} it was much more likely that some other simple two-digit addition formula would have popped into his head to consider. Furthermore, given that two-digit addition formulas only occupy a tiny portion of *all* simple arithmetical formulas, the vast majority of Peter's arithmetical judgments throughout the afternoon *were not* judgments on 2-digit addition formulas.

Once again, it seems as if t_{c3} and t_{m3} fail to generate knowledge even though their resultant judgments are true. Interestingly, given the counterfactually nearby possibility space for these tokens, it seems as if t_{c3} is reliable with respect to C_p, C_n and C_b, and that t_{m3} is reliable with respect to M_p, M_n and M_b. In both of these cases, the absence of knowledge appears to stem from a lack of reliability with respect to a content-evidence pair with an *intermediate* degree of specificity. These types, for t_{c3} and t_{m3} respectively, are as follows:

C_I [judging an object to have a particular shade of orange, on the basis of visual experiences that represent a particular color shade for that object]

M_I [judging the value of a simple two-digit addition formula, on the basis of an intuitive seeming representing a value for that given formula]

Upon reflection, there appears to be a general recipe for creating cases in which Susie fails to have knowledge: insofar as she is disposed to make mostly incorrect judgments with respect to *any* content-evidence pair *with any degree of specificity* that falls between C_p and C_b, her token will not generate knowledge. The same can be said for Peter, M_p, and M_b.

We see here that Conee and Feldman's *single-type* understanding of the generality problem radically mis-describes the nature of knowledge-enabling reliability. Whether a token generates knowledge seems to depend on reliability measurements taken across many content-evidence pairs.

But how should we conceive of the varying degrees of descriptive specificity that characterize all of a token's relevant types? How exactly do the reliability scores for *all* of these types come together to determine whether a token generates knowledge? In what follows, I attempt to make headway on these questions.

4. Similarity Relations and Multiple Process Types

As I argue below, there is an elegant structure for analyzing all of a token's relevant types. As it turns out, this structure closely mirrors the organization of knowledge-determining counterfactual reference classes.

4.1 Counterfactual Distance and Reference Classes

Consider Susie once again. As we saw from COLOR 1, Susie's token ascriptions of orange 47 can generate knowledge only if they are counterfactually reliable with respect to the broader type C_b. The RC account of reference classes captures token t_{c1}'s lack of knowledge-enabling reliability in the following way: there is a low truth ratio measured across the class of all possible belief-forming events counterfactually close to t_{c1} in which Susie undergoes C_b processes. As I discussed above, formulating reference classes counterfactually allows us to elegantly characterize the way in which a subject's belief-forming dispositions determine whether their beliefs count as knowledge.

According to contemporary counterfactual semantics, possibilities that are counterfactually "closer together" have a higher degree of similarity to each other, whereas possibilities that are counterfactually "further apart" are less similar.[25] In other words, degrees of counterfactual distance are simply degrees of *difference* between two or more possibilities. Of course, there are many different similarity relations, i.e., many distinct ways in which things can be similar or different from each other. For simplicity, let "K" denote the specific similarity relation that grounds the counterfactual distance measurements constitutive of a token's knowledge-determining reference classes.

[25] The notion of counterfactual "distance" as being determined by a similarity/ordering relation of possible worlds plays a central role in both Lewis (1973) and Stalnaker (1968) semantics for counterfactual sentences.

RC incorporates the following reasonable idea: whether a token has knowledge-enabling reliability does *not* depend on what happens in possible scenarios that are *very different* from the actual token case. For example, there are possible belief-forming events in which Susie, unbeknownst to her, is a brain-in-a-vat. In these BIV possibilities, all of Susie's C_b judgments are false, given that they are all based on hallucinatory experiences. However, it doesn't seem like these false judgments in the BIV possibilities drive down the knowledge-determining C_b truth ratios for the tokens in COLOR 1-3. This suggests that the BIV belief-forming possibilities are *not* members of the knowledge-determining reference classes belonging to t_{c1}, t_{c2}, and t_{c3}. The BIV possibilities are simply too different from t_{c1}, t_{c2}, and t_{c3} to matter for knowledge-enabling reliability. We can plausibly infer that a token's reference classes are only comprised of possibilities that all lie within some maximum counterfactual distance D_m from the token itself—relative to similarity relation K. In the case discussed above, it's clear that the BIV possibilities lie outside of D_m measured from tokens t_{c1}, $t_{c2,}$ and t_{c3}.

As RC describes, there is a distinct reference class for each relevant type T belonging to a given process token t. As a result, each reference class belonging to a token can be geometrically represented as a sphere of possibility space (see Figure 1). For any token t and relevant type T, t is at the center of the sphere, and the radius is counterfactual distance measurement D_m. The volume within the sphere is comprised by all belief-forming events within D_m in which the subject undergoes T-exemplifying processes.

Figure 1.

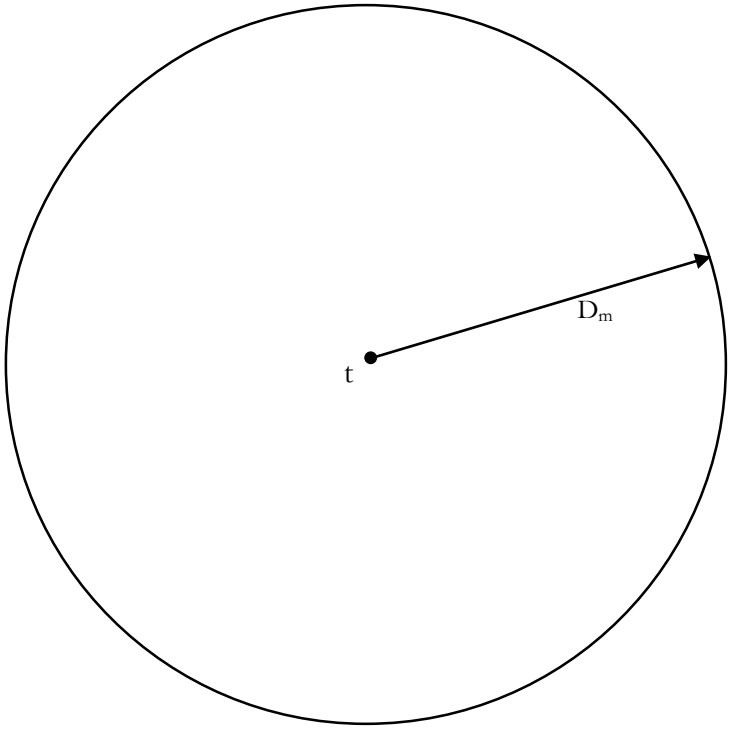

4.2 Conceptual Difference and Multiple Process Types

As it turns out, the counterfactual structure for reference classes provides an apt model for analyzing a token's many relevant *types* with each of their differing degrees of descriptive specificity. To begin, consider COLOR 1-3. In all three of these cases, the maximally precise content-evidence pair used in the token case is C_p. Relative to C_p, C_b has a lesser degree of

specificity than C_n. After all, the extension of C_n is a rather small set of precise content-evidence pairs including,

> [forming a judgment on whether a given object x is orange 47, on the basis of a visual experience that represents a given object x to be orange 47]
>
> [forming a judgment on whether a given object y is orange 47, on the basis of a visual experience that represents a given object y to be orange 47]
>
> [forming a judgment on whether a given object z is orange 47, on the basis of a visual experience that represents a given object z to be orange 47]
>
> …

Importantly, each member of C_n's extension is, conceptually, very similar to C_p. On the other hand, the extension of C_b is a vast set of precise content-evidence pairs, including everything from [judging whether a given object x is brown 11 on the basis of a visual experience that represents the object x to be brown 11] to [judging whether a given object u is yellow 12, on the basis of a visual experience that represents the object u to be yellow 12], and many more. While every member of C_n's extension is, conceptually, very similar to C_p, many members of C_b's extension are significantly less similar to C_p.

This observation suggests that we can define types like C_b and C_n—and thereby capture their varying degrees of specificity—in terms of a similarity relationship to C_p. Whereas reference classes are defined in terms of a *counterfactual distance* measurement taken from the token event t, relevant content-evidence pairs can be defined in terms of a *conceptual difference* measurement taken from the token's precise content-evidence pair. For simplicity, let "L" denote the similarity relation that grounds the conceptual difference measurements constitutive of a token's relevant content-evidence pairs. The degree of conceptual difference from C_p that characterizes C_b is *greater than* the degree of conceptual difference from C_p that characterizes C_n. More generally, relevant types can be analyzed in the following way:

> Relevant Type Organization (RTO)
>
> For every token t and precise content-evidence pair CE_t exemplified by t, each of t's relevant content-evidence pairs CE_i can be analyzed in terms of a particular conceptual difference measurement E_i taken from CE_t, according to similarity relation L.

Furthermore, relevant content-evidence pairs can be represented as spheres as well. For a given token t, precise content-evidence pair CE_t, and relevant type CE_i, the sphere corresponding to CE_i has CE_t as the center, while the radius is the particular degree of conceptual difference E_i that characterizes CE_i. The extension of CE_i comprises the volume within the sphere.

While it's clear that the narrower C_n's extension is a subset of the broader C_b's extension, representing these types as spheres shows C_n's extension to be a *concentric sphere* of the larger C_b's extension relative to center point C_p. In Figure 2, the grey sphere represents C_b and the horizontal-lined sphere represents C_n. For the sake of argument, let's assume that the shade orange 47 lies right in the middle of the orange section of the color spectrum. With this assumption, we can model type C_I with the dotted sphere, given that it has a degree of

specificity (relative to C_p) that falls in between that of C_n and C_b. Likewise, for MATH 1-3, types M_n, M_I, and M_b can be represented as concentric spheres of each other relative to center point M_p.

As we saw with reference classes, there's good reason to believe that truth ratios across belief-forming possibilities *very different* from the token are irrelevant to whether that token has knowledge-enabling reliability. Similarly, for relevant types, it seems that truth-ratios taken across instances of content-evidence pairs that are vastly different from the token's precise content-evidence pair have no bearing on whether that token generates knowledge. To see this, consider the following case:

> MATH COLOR INVERSION
>
> Suppose that Peter's accident didn't alter his arithmetical intuition abilities at all, but instead (unbeknownst to him) made him red/green spectrum inverted. Now, red things visually appear green to him, and vice versa. As a result, he's unreliable with respect to the content-evidence pair RT: [judging whether an object x is red on the basis of a visual experience that represents x to be red]. Suppose that, as he is practicing arithmetic, he undergoes token t_{m4}, in which he intuits 37+14=51. As it turns out, t_{m4} occurs while Peter is standing in the middle of a red and green playground. Quite often, Peter takes a break from arithmetic to observe his environment.

Given Peter's location at the time t_{m4} occurs, many of the possible judgments that are counterfactually close to t_{m4} are RT judgments. Given Peter's accident, these counterfactually nearby RT judgments will be mostly false. However, these counterfactually nearby false judgments don't seem to have any bearing on whether token t_{m4} generates knowledge. Plausibly, possible instances of RT are not included in any of t_{m4}'s knowledge determining truth-ratio measurements. RT simply seems *too* different from t_{m4}'s precise content-evidence pair (M_p) to matter for any of t_{m4}'s knowledge-determining truth-ratio measurements. We can rather easily multiply cases like this, where it seems that counterfactually nearby judgments featuring content-evidence pairs that are vastly different than the token's precise content evidence pair won't have any impact on whether that token has knowledge-enabling reliability. This suggests that there is a finite, maximum degree of conceptual difference E_m that characterizes a token's broadest relevant content-evidence pair description. In the case discussed above, it's clear that instances of RT are *not* included in the extension belonging to the E_m relevant type for t_{m4}.

In Figure 2, E_m represents the degree of conceptual difference that characterizes the broadest relevant type belonging to Susie's token processes in COLOR 1-3.

Figure 2.

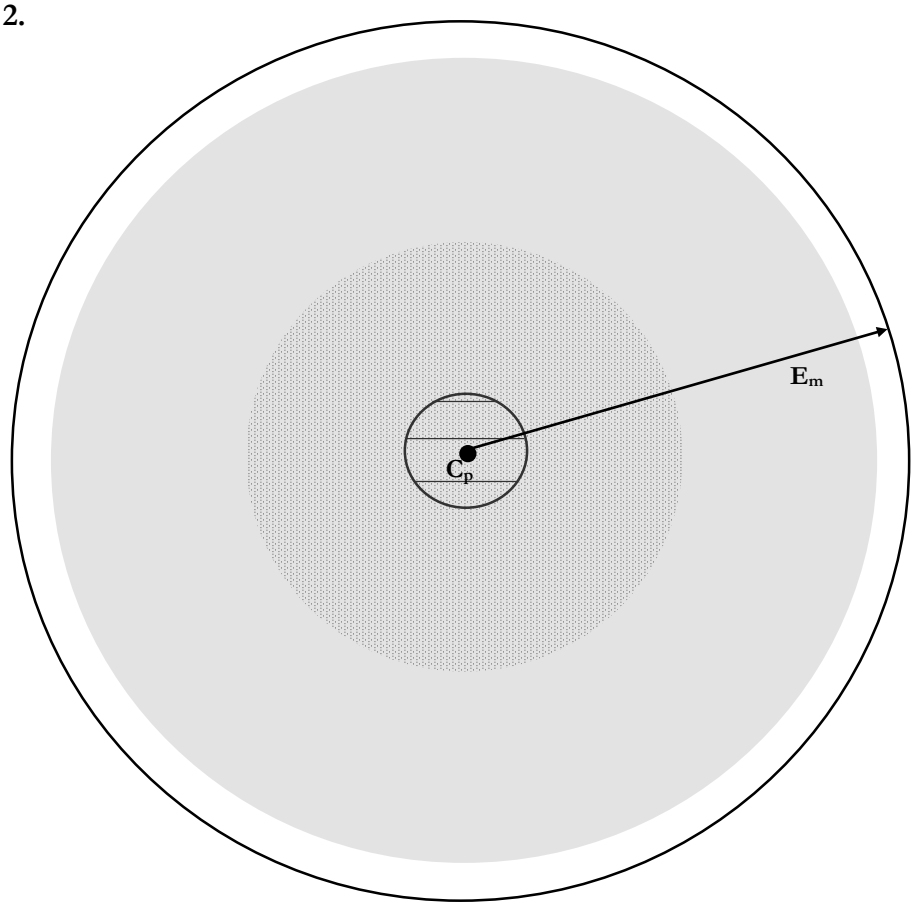

As we saw, it's rather easy to multiply cases like COLOR 1-3 and MATH 1-3 to see that, in order to generate knowledge, tokens must be counterfactually reliable with respect to *numerous* content-evidence pairs, each with varying degrees of descriptive specificity. Furthermore, invoking degrees of conceptual difference allows us to formulate an account of knowledge-enabling reliability that picks out *all* of a token's relevant types. I call this account multi-type evidential reliabilism (MTE):

MTE For any belief-forming process token t and subject S, where CE_t is the precise content-evidence pair description exemplified by t, t generates knowledge only if,

For *each* concentric content-evidence pair description CE_i, characterized by a particular degree of conceptual difference E_i that is $\leq E_m$ (measured from CE_t),

CE_i has a sufficiently high truth-ratio with respect to the reference class of possibilities (in which S undergoes processes that instantiate CE_i) that fall within counterfactual distance D_m from t.[26]

[26] Most importantly, I hold that MTE is one of the necessary conditions on knowledge. Moreover, I think MTE crystalizes the sort of multi-type *structure* that characterizes knowledge-enabling reliability. However, for all I've said here, there *could* be additional necessary conditions on knowledge that pertain to knowledge -

The "concentric-sphere" structure of relevant types is a key aspect of MTE.[27] As Figure 2 illustrates, types C_b, C_I, and C_n are concentric sphere types relative to center point C_p. However, there are some types exemplified by tokens t_{c1}, t_{c2}, and t_{c3} that are *not* concentric spheres relative to C_p. For example, consider the following type:

> C_V [forming a judgment *either* on whether an object is orange 47 *or* on whether an object is any particular shade of violet, on the basis of a visual experience representing a color shade for that object][28]

Importantly, violet shades occupy just one part of the color spectrum that's rather distant from where the orange shades stand on the color spectrum. With this in mind, we can visually represent the relationship between the broad type C_b and type C_V with Figure 3. The grey space represents C_b and the checkered space represents C_V. C_p lies in the center, and the radii measured from C_p corresponds to degrees of conceptual difference relative to relation L. Given the disjunctive structure of C_V, its spatial representation in Figure 3 occupies the overlap between [forming a judgment on whether an object is orange 47, on the basis of a visual experience representing a color shade for that object] and [forming a judgment on whether an object is any particular shade of violet, on the basis of a visual experience representing a color shade for that object]. Unlike C_b, C_V does *not* have a concentric spherical structure relative to center point C_p.

---

enabling reliability. For example, Duncan Pritchard has recently suggested that, in order to generate knowledge, tokens must manifest reliability across multiple reference classes each characterized by *different counterfactual distances* from the actual token case (2012: 179-80). MTE could be augmented to incorporate such a reliability requirement across numerous counterfactual distances, although exploring this option is beyond the scope of this paper. Furthermore, one might think that there are *multiple* similarity relations in addition to L that determine a token's relevant types. On this picture, a token would have multiple concentric sphere type structures, each corresponding to a different similarity relation.

[27] See fn. 20. More recently, Duncan Pritchard has clarified that a belief that p is safe if and only if there is a sufficiently high truth-ratios measured across nearby possible judgments regarding p made "on the same basis" as the actual belief itself (2012: 176, 179). If MTE is correct, then we can reasonably read Pritchard's safety condition on knowledge as constituting *one part* of the reliability condition on knowledge. In particular, Pritchard's safety condition expresses the reliability requirement with respect to just the token's precise content-evidence pair (CE_t).

[28] Thanks to an anonymous referee for raising the example of C_V and encouraging me to clarify the difference between relevant types according to MTE and irrelevant types like C_V.

Figure 3.

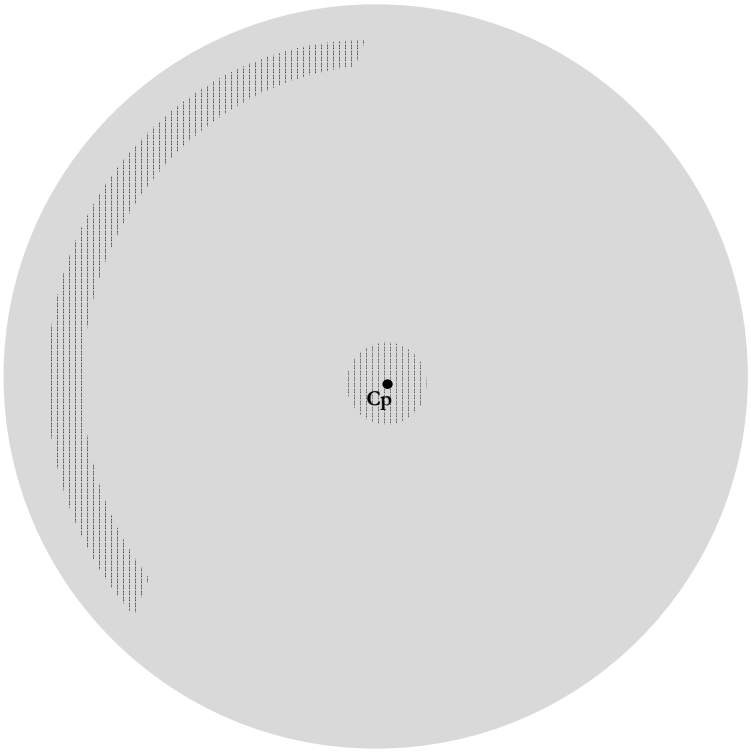

Upon considering cases like COLOR 1-3, it seems that any token process with precise content-evidence pair C_p must be reliable with respect to C_n, C_I, and C_b in order to generate knowledge. On the contrary, it does *not* seem that such tokens must also be reliable with respect to non-concentric types like C_V in order to generate knowledge. To see this, consider the following scenario:

> PLANET X
>
> Much like Susie, Jim likes to visually identify the color shades of objects around him. Unlike Susie, Jim lives on Planet X, which is much like Earth except for one key difference: while there are some violet-colored objects scattered throughout Planet X, there are *very few* of them. Also, unbeknownst to Jim, he's acquired a cognitive abnormality that gives him a stable disposition to visually misrepresent the specific shades of any violet-colored object. His ability to correctly represent the shades of any other color are left untouched by the abnormality. One day, Jim looks at an orange 47 object, visually represents it to be orange 47, and on that basis comes to believe the object to be orange 47 in token process t_{J1}. Around the time t_{J1} occurs, there are a few violet objects in Jim's vicinity, making it roughly just as likely—with respect to nearby possibility space—that Jim would have looked at a violet object as it was that he would have looked at an orange 47 object. That said, given that Jim lives on Planet X, color-shade judgments about violet objects make up an extremely tiny percentage of *all* of Jim's counterfactually nearby (visually-based) color-shade judgments at the time and place in which t_{J1} occurs.

In PLANET X, token t_{J1} is counterfactually unreliable with respect to type C_V. However, it does not seem that unreliability with respect to C_V undermines Jim's knowledge in this case. All else being equal, it seems that t_{J1} manifests knowledge-enabling reliability given that it is counterfactually reliable with respect to C_p, C_n, C_I, C_b, etc. MTE accommodates this intuition

given that it only requires that a token manifest reliability with respect to *concentric* content-evidence pairs in order to generate knowledge.[29]

Given the explanatory power of MTE across a variety of cases, we have good reason to believe that MTE successfully meets the specificity challenge. In what follows, I return to Conee and Feldman's generality problem objection and explore the current state of this argument against reliabilism.

5. Responding to the Generality Problem Objection

The generality problem objection claims that we have good reason to reject a reliability condition on knowledge because the generality problem looks *unsolvable*. However, I think that MTE makes sufficient progress towards solving the generality problem so as to undermine the generality problem objection. On this point, it's instructive to take stock of the theoretical progress made by MTE.

To begin, as I argued above, the long track record of failure to solve the generality problem—and the apparent intractability of the specificity challenge—stems from the faulty assumption that knowledge-enabling reliability is determined by the truth ratio of just a *single* type. MTE rejects this single-type assumption and meets the specificity challenge by describing knowledge-enabling reliability with a concentric multi-type structure. According to MTE, relevant types are content-evidence pairs, and every relevant type that belongs to a given token has a different degree of descriptive specificity. These differing degrees of specificity are defined by a similarity relation and measurements of conceptual difference. Furthermore, the degrees of specificity that characterize a token's relevant types are determined by degrees of similarity *relative to the precise content-evidence pair* instantiated by the token. In this way, MTE incorporates Comesaña's insight that the token's precise content-evidence pair plays a central role in determining a resultant belief's epistemic status. According to MTE, a token's precise content-evidence pair is something akin to an "anchor-point" for determining all of the other relevant types. As we saw, MTE's unique approach to explicating knowledge-enabling reliability accommodates our intuitions on a wide variety of cases (e.g., MATH 1-3, COLOR 1-3)—cases that haven't been addressed by other extant responses to the generality problem.

That said, for all of its theoretical contributions, it's important to note that MTE remains somewhat schematic in its current form. MTE analyzes knowledge-enabling reliability by invoking these four key notions:

---

[29] That being said, according to MTE it is still *possible* for someone's visual incompetence at representing violet shades to undermine his knowledge in cases where one visually ascribes orange 47 to some object. It all depends on the distribution of possible belief-forming events that are counterfactually close to the token process. Suppose that, rather than living on Planet X, Jim lives on Planet Y where 80% of all things are colored some shade of violet. In this scenario, when Jim undergoes process token t_{J2} in which he visually identifies an object as being orange 47, he is surrounded by mostly violet-shaded objects. Here, it does seem that Jim's inability to visually represent violet shades undermines the knowledge-enabling reliability of t_{J2}. MTE can straightforwardly accommodate this result, because t_{J2} is counterfactually unreliable with respect to the concentric type C_b in this altered scenario. Relative to t_{J2}, the majority of counterfactually nearby color shade judgments would be judgments *about violet objects* (given that t_{J2} occurs on Planet Y). But given that Jim is disposed to misrepresent violet shades, the truth ratio across all counterfactually nearby C_b judgments would be low.

(a) Evidence
(b) Basing
(c) Similarity relations K and L
(d) Particular degrees of difference D_m and E_m

In this essay, I have *not* taken it upon myself to thoroughly defend any informative account of the evidential basing relationship. Secondly, while I presented cases to support the idea that type relevance is determined by similarity relations and degrees of difference, I have *not* offered or defended informative accounts of these similarity relations and degrees of difference. Hence, there is additional work to be done in solving the generality problem, as any fully satisfying and informative theory of type-relevance would further explicate notions (a)-(d).[30] Reasonably, developing informative accounts of (a)-(d) is a task that lies outside the scope of this paper.[31]

Nevertheless, as it pertains to the generality problem *objection,* I contend that MTE successfully undermines this argument against reliabilism. MTE illuminates the nature of knowledge-enabling reliability to the point where the generality problem no longer seems *unsolvable*. MTE illustrates the multi-type structure that any fully informative account of knowledge-enabling reliability must have while highlighting the key concepts in need of further exploration for future work on the generality problem.

Now of course, MTE—in its current form—wouldn't constitute an adequate rebuttal to the generality problem objection if it explicated knowledge-enabling reliability only by invoking notions that are even more mysterious and opaque than the concept *relevant type* itself. As Jonathan Matheson notes, such an account merely "swaps out one problem for an equally bad problem" (2015: 467). By way of response, there's good reason to believe that (a)-(d) are *the kinds concepts* that are both familiar to us in everyday life, and commonplace in philosophical

[30] Thanks to an anonymous referee for identifying the current limitations of MTE and work that remains to be done in order to solve the generality problem.

[31] As Korcz illustrates, there already exists a vast and highly contentious literature on the nature of the epistemic basing relationship (2015). Furthermore, at this stage of the inquiry, one could reasonably foresee various approaches that might be taken to further explicate concepts (c) and (d). Take similarity relation L for instance. The contemporary work on similarity/ordering relations in the counterfactual semantics literature might provide insights into the nature of L as well. David Lewis argues that counterfactual similarity relations are determined by features of the context in which the counterfactual sentence is uttered (1979: 465). Analogously, one might argue that relation L functions in a similar way, thus making for an interesting sort of contextualism about knowledge attributions that arises due to the contextual nature of process type relevance. Another option worth exploring would be to analyze relation L in terms of *understanding*. Plausibly, in order to *base* some belief that p_1 on evidence e_1, the subject must possess some sort of understanding of the relationship between p_1 and e_1. Let U_1 denote this cognitive state of understanding that corresponds to content-evidence pair $[p_1,e_1]$. Presumably, there's a unique cognitive state of understanding that corresponds to each distinct content-evidence pair. For example, let U_2 denote the state of understanding corresponding to $[p_2, e_2]$. Perhaps the epistemically relevant degree of similarity (according to L) that holds between content-evidence pairs $[p_1,e_1]$ and $[p_2,e_2]$ is determined by the degree of similarity between cognitive states U_1 and U_2. The challenge for this approach lies in explicating the kind of understanding involved in epistemic basing. Lastly, after further exploration, one might uncover good reason to conclude that L is conceptually primitive, admitting of no further informative analysis or explanation. Certainly, this is a result that we're in no position to rule out at this point in time. Of course, some other approach to explicating L may emerge—these are just preliminary suggestions. But what's important to notice at this point, regarding the generality problem *objection,* is that we have *no reason to doubt* that further theoretical progress on notions like L *can* occur.

theorizing. Reasonably, explicating type-relevance with these sorts of concepts is sufficient for salvaging the plausibility of reliabilism from the generality problem objection.

To begin, Comesaña points out that "any adequate epistemological theory needs to appeal, either implicitly or explicitly, to the notion of a belief's being based on certain evidence" (2006: 33). Indeed, notions (a) and (b) play central roles in Conee and Feldman's own favored theory of epistemic justification: evidentialism. According to evidentialism, believing p is propositionally justified for subject S if and only if believing p *fits* the evidence E that S possesses (Conee and Feldman 1985: 15). Furthermore, on evidentialism, a subject S's belief that p is doxastically justified if and only if believing p is propositionally justified for S, and S bases her belief that p on her evidence E (1985: 34).[32] Because of this, Comesaña notes that invoking notions (a) and (b) to answer the generality problem "should be accepted not only by Conee and Feldman, but by anyone who thinks that an epistemological theory is incomplete without an appeal to the basing relation" (38).

In contrast, there's no reason to think that non-reliabilists would be committed to the specific notions (c) and (d). However, I think that non-reliabilists of many varieties are committed to notions that are quite similar to (c) and (d) insofar as these theorists countenance dispositional or counterfactual concepts. As I explained above, MTE describes *knowledge* as a dispositional concept. As many cases—including COLOR 1-3 and MATH 1-3—suggest, whether a token process generates knowledge crucially depends on the cognitive dispositions of the subject undergoing that process. According to reliabilism, at least some of these knowledge-determining dispositions can be described as particular sorts of truth-ratios. Moreover, we've seen that counterfactual notions, like similarity relations and particular degrees of difference, provide a reasonable and elegant way of characterizing these dispositional profiles.

The use of dispositional concepts is pervasive in everyday thought and speech. Consider virtue concepts, like *honesty*. Presumably, whether someone possesses honesty in a given situation X crucially depends on what she is disposed to do in situation X. It's rather easy to think of many other commonplace dispositional concepts like this: *being courageous, being a skilled archer, being physically coordinated,* etc. Furthermore, it makes sense to view these dispositional concepts as having truth conditions framed in terms of similarity relations and particular degrees of difference. For example, whether S is honest in situation X seems to (at least) depend on what S would do in a variety of possible situations that are sufficiently similar to X. Plausibly, what counts as "sufficiently similar to X" is determined by a similarity relation of some sort and some maximum degree of difference from X.

More generally, humans display a rather significant aptitude for correctly deploying counterfactual concepts across a variety of situations. In numerous circumstances, humans can proficiently communicate and reflect on how things *would* have been if events *had* gone a bit differently. Given contemporary counterfactual semantics, *any* counterfactual truth is determined by a similarity relation and a particular degree of difference.

For our purposes, it's important to note that humans can competently deploy these counterfactual and dispositional concepts without explicitly possessing any sort of informative account of either their constitutive similarity relations or maximum (relevant) degrees of difference. Reasonably, the vast majority of humans grasp concepts like *honesty* in virtue of having a tacit sensitivity to the workings of *honesty's* constitutive similarity relation and relevant

[32] Interestingly enough, Comesaña, in his presentation of the well-founded solution, seems to assume the same account of evidence as Conee and Feldman: evidence consists of mental states (Conee and Feldman 2008: 84-88).

maximum degree of difference. Likewise, I think that most humans have precisely this sort of tacit sensitivity to (c) and (d), which allows us to correctly deploy the knowledge concept in a variety of situations. It appears that MTE meets the specificity challenge by describing knowledge-enabling reliability with the sorts of notions that most humans grasp.

At this point, those sympathetic to the generality problem objection might insist that defenders of MTE owe us an informative account of its central dispositional concepts (c) and (d), and that without this supplementary account, MTE fails to adequately undermine the generality problem objection. But such a demand would make the generality problem objection implausibly too strong. In particular, this demand would seemingly overgeneralize to undermine *any* philosophical theory that invokes key dispositional notions without analyzing them further.[33]

For example, this demand for theories to include supplementary accounts of any key dispositional concepts would appear to undermine Conee and Feldman's own evidentialist theory of propositional justification as well. To see this, consider how evidentialists invoke the notion of *evidence possession.* One might wonder what it takes to possess some piece of evidence. Lately, evidentialists have defended a dispositional account of evidence possession. According to a dispositional account of evidence possession, whether a subject S possesses e as evidence for p crucially depends on the thoughts that S is disposed to have regarding p and e.[34] In a previous discussion of the generality problem, I note that evidentialists haven't provided anything like an informative account of the dispositions that ground one's possession of evidence (2017: 1952-3).[35] There, I conclude that, *if* the absence of a supplemental account of dispositional notions undermines reliabilism, then it would also undermine evidentialism. But, upon reflection, evidentialism *doesn't* seem to be undermined simply by its lack of such supplemental accounts. By parity of reasoning, we also shouldn't think that reliabilism coupled with MTE loses its plausibility merely because it lacks a supplementary account of its dispositional notions (c) and (d).

In sum, MTE undermines the generality problem objection by showing how the generality problem can be solved. Moreover, MTE describes type-relevance in a way that meets the specificity challenge by invoking the sorts of notions that are commonplace in philosophical theorizing. Lastly, MTE identifies which key notions still require further exploration in the task of giving a fully informative theory of knowledge-enabling reliability.

---

[33] For example, many philosophers, including Price (1969) and Audi (1972), think that *belief* is a dispositional concept. On this view, whether one believes that p in situation W depends on whether she has dispositions to think or act in certain ways in W. The dispositional theory of belief has broad appeal and seems reasonable as stated. To the point in question, it would be odd if the reasonability of the dispositional theory of *belief* were in some way undermined simply due to an absence of further theoretical work to explicate *belief's* corresponding similarity relation and relevant maximum degree of difference.

[34] Evidentialist Kevin McCain has recently offered a substantive and compelling defense of a dispositional account of evidence possession (2014: 31-55). I defend a dispositional account of evidence possession as well (2017: 1949-1950). According to both McCain and myself, if evidence possession (at a given time t) only depends on one's *occurrent* mental states, then evidentialism won't be able to account for the vast majority of the justified beliefs that we plausibly have.

[35] Additionally, I argue that there's no reason to view these sorts of dispositions as importantly different or less complex than the dispositions invoked by reliabilism (1951-3).

Conclusion

In this essay, I've defended a new multi-type approach to solving the generality problem that explicates the reliability condition on knowledge. Where single-type theories have failed, MTE successfully meets the specificity challenge by describing the *multiple* relevant types that determine whether a given token generates knowledge. All the while, MTE retains Comesaña's key insight that knowledge-enabling reliability crucially depends on a token's precise content-evidence pair. According to MTE, all of a token's relevant types are defined in terms of varying degrees of conceptual difference measured from the token's precise content-evidence pair. Moreover, MTE successfully undermines the famed generality problem objection to reliabilism by meeting the specificity challenge while only invoking familiar concepts that are relatively common throughout philosophy in general.[36]

Resources

Adler, Jonathan and Michael Levin. (2002), "Is the Generality Problem too General?" *Philosophy and Phenomenological Research*. 65.1: 87-97.

Alston, William. (1995), "How to Think about Reliability." *Philosophical Topics*. 23.1: 1-29.

Audi, Robert. (1972), "The concept of believing." *Personalist*. 53: 43–62.

Becker, Kelly. (2008), "Epistemic Luck and the Generality Problem." *Philosophical Studies*. 139.3: 353-366.

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.

Brueckner, Anthony and Christopher T. Buford. (2013), "Becker on Epistemic Luck." *Philosophical Studies*. 163: 171-175.

Conee, Earl and Richard Feldman. (1985), "Evidentialism." *Philosophical Studies*. 48:15-34.

——— (1998), "The Generality Problem for Reliabilism." *Philosophical Studies*. 89.1:1-29.

——— (2002), "Typing Problems." *Philosophical and Phenomenological Research*. 65.1:98-105.

——— (2008), "Evidence." *Epistemology: New Essays*. Ed. by Quentin Smith. Oxford: OUP. 83-104.

[36] Much thanks to Ted Warfield, Blake Roeber, Tom Senor, Liz Jackson, Ting Cho Lau, Andrew Moon, and Peter Finocchiaro for helpful conversations and comments at various stages of drafting this paper.

Comesaña, Juan. (2002), "The Diagonal and the Demon." *Philosophical Studies*. 110: 249–266.

——— (2006), "A Well-Founded Solution to the Generality Problem." *Philosophical Studies*. 129.1: 27-47.

———(2010), "Evidentialist Reliabilism." *Noûs*. 44.4:571-600.

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: D. Reidel Publishing. 1-23.

——— (1986), *Epistemology and Cognition*. Harvard: Harvard University Press.

——— (1988), "Strong and Weak Justification." *Philosophical Perspectives*. 2: 51-69

Greco, John. (2010), *Achieving Knowledge*. Cambridge: CUP.

Heller, Mark. (1995), "The Simple Solution to the Generality Problem." *Noûs*. 29.4:501-515.

Kappel, Klemens. (2006), "A Diagnosis and Resolution to the Generality Problem." *Philosophical Studies*. 127: 525-60.

Korcz, Keith Allen. (2015) "The Epistemic Basing Relation." *The Stanford Encyclopedia of Philosophy* (Fall 2015 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2015/entries/basing-epistemic>.

Lepock, Christopher. (2009), "How to Make the Generality Problem Work for You." *Acta Analytica*. 24: 275-286.

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

Lewis, David. (1973), *Counterfactuals*. Cambridge: Harvard University Press.

——— (1979), "Counterfactual Dependence and Time's Arrow." *Noûs*. 13.4: 455-476.

Matheson, Jonathan D. (2015), "Is there a Well-Founded Solution to the Generality Problem?" *Philosophical Studies*. 172: 459-468.

McCain, Kevin. (2014), *Evidentialism and Epistemic Justification*. New York: Routledge.

Olsson, E. J. (2016), "A Naturalistic Approach to the Generality Problem." *Goldman and His Critics*. Ed. by B. P. McLaughlin and H. Kornblith. Chichester: Wiley-Blackwell. 178-199.

Price, H.H. (1969), *Belief*. London: George Allen and Unwin.

Pritchard, Duncan. (2005), *Epistemic Luck*. Oxford: OUP.

——— (2012), "A Defense of Modest Anti-Luck Epistemology." *The Sensitivity Principle in Epistemology*. Ed. by Kelly Becker and Tim Black. Cambridge: CUP. 173-192.

Reichenbach, H. (1949). *The Theory of Probability*. University of California Press.

Schmitt, Frederick F. (1992), *Knowledge and Belief*. London: Routledge.

Sosa, Ernest. (1991), *Knowledge in Perspective*. Cambridge: CUP.

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

Tolly, Jeffrey. (2017), "A Defense of Parrying Responses to the Generality Problem." *Philosophical Studies*. 174.8: 1935-1957.

——— (forthcoming), "Swampman: a Dilemma for Proper Functionalism." *Synthese*.

Wallbridge, Kevin. (2016), "Solving the Current Generality Problem." *Logos and Episteme*. 7.3: 345-50.

Wallis, Charles. (1994), "Truth-Ratios, Process, Task, and Knowledge." *Synthese*. 98: 243-269.

Wunderlich, Mark. (2002), "Vector Reliability: A New Approach to Epistemic Justification." *Synthese*. 136: 237-62.