

Complexly-Based Beliefs and the Generality Problem for Reliabilism¹

Max Baker-Hytech

Wycliffe Hall, Oxford University

Introduction

This paper argues that certain cases involving what I shall term “complexly-based belief,” where a belief is formed via complex inference to the best explanation, pose a serious difficulty for reliabilist theories of epistemic justification or warrant. Many of our most important beliefs appear to be of this character. The problem, in short, is that in such cases we cannot identify *any* belief-forming process type which is such as to yield an intuitively correct verdict on the epistemic status of the agent’s belief. If this is correct, then no proposed solution to the generality problem can succeed.

The generality problem for process reliabilism and modal reliabilism

In broad terms, the generality problem is the problem of finding a principle or rule for determining which general type of belief-forming process should be regarded as the relevant one for assessing the reliability of a given belief. Process reliabilism can be understood in the following way, where “warrant” denotes that property enough of which turns a true belief into an item of knowledge:²

(PR) S’s belief *p* is warranted to the extent that the causal process which produced that belief is reliable.³

The reason that the generality problem arises for (PR) is as follows. Any given belief is the result of a particular causal sequence or process. For instance, Jim’s belief that *the dinner is ready* is the end result of the following causal sequence: first, Jim put the dinner in the oven and set the timer, which caused him to have the belief that *the dinner will be ready when the timer goes off*; then later, the timer went off, which in conjunction with Jim’s other beliefs, caused Jim to have the belief that *the dinner is ready*. This is a rough and ready description of the causal process that led to Jim’s belief. We ask the question “is the process reliable?” If by “the process” we mean the sequence of

¹ I am grateful to C’Zar Bernstein, Ben Page, Logan Gage, and Alex Plato for very helpful comments on earlier versions of this paper.

² Here I am following Alvin Plantinga’s usage in *Warrant: The Current Debate* (Oxford: Oxford University Press, 1993).

³ Alvin Goldman is the foremost advocate of an account along these lines. See his “What Is Justified Belief?,” in *Justification and Knowledge*, ed. George S. Pappas (Dordrecht: D. Reidel, 1979). See also Kelly Becker, *Epistemology Modalized* (London: Routledge, 2007); John Greco, *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity* (Cambridge: Cambridge University Press, 2010). These latter two accounts of warrant include a process reliability condition, but more besides.

events that caused Jim's belief, then it isn't meaningful to ask whether "the process" is reliable, because reliability is a tendency, a matter of having a high success rate across a suitably defined reference class. If "the process" denotes the particular sequence of events that caused Jim's belief, then we can't sensibly ask whether it is reliable, because there isn't any sense in which this particular sequence of events produces a true belief in most cases. There is only one case for us to consider.

Hence why, when we ask whether "the process" which produced Jim's belief was reliable, we must really be asking about the reliability of the general *type* of process instantiated by the particular causal sequence that led to Jim's belief. What, then, is the general type that is instantiated here? Immediately one realizes that there are many candidates: *inference from perception and memory*; *hearing the oven timer go off*; *perceiving informative signals from domestic appliances*; and so on. All of these general types are instantiated in the token causal process that produced Jim's belief. The generality problem, then, is the problem of devising a principled way of selecting the uniquely relevant process type for any given token belief.

The generality problem also arises for modal reliabilist accounts of warrant.⁴ A modal reliabilist account of warrant appeals to the notion that the agent wouldn't have formed a false belief in certain counterfactual circumstances. The sensitivity condition and the safety condition are both modal reliability conditions. They differ in specifying different counterfactual circumstances in which error needs to be avoided by the agent. The sensitivity condition requires that the agent would not have believed *p* in the nearest possible world(s) in which *p* is false.⁵ Depending on the proposition, the nearest world in which *p* is false may be very distant from the actual world. For example, suppose that in the actual world we are not brains in vats being subjected to massive deception; reality is roughly as we suppose it to be. The nearest world in which this proposition is false—i.e. the nearest world in which we *are* brains in vats being subjected to massive deception—is a very distant world. Naturally, given that in such a world we are being subjected to massive deception, we fail to be aware of the truth that we are being so subjected. Hence, my belief that I am not a brain in a vat fails to satisfy the sensitivity condition. The safety condition for knowledge merely requires that the agent avoids false belief about *p* in nearby possible worlds, regardless of the content of *p*.⁶ Hence, given that the nearest world in which I am a brain in a vat who doesn't realize he is a brain in a vat is a very distant world, my true belief that I am not a brain in a vat can satisfy the safety condition.⁷

⁴ The generality problem for modal reliabilism receives less discussion than the generality problem for process reliabilism, but the problem as it pertains to the sensitivity condition has been discussed by Mark Alfano, "Sensitivity Theory and the Individuation of Belief-Formation Methods," *Erkenntnis* 70 (2009). Timothy Williamson acknowledges that "The generality problem arises for my safety account because, like process reliabilism, it faces the question 'How similar must counterfactual processes [bases] be to count towards reliability/safety?'" ("Reply to Alvin Goldman," in *Williamson on Knowledge*, ed. Patrick Greenough and Duncan Pritchard (Oxford: Oxford University Press, 2009), 100.)

⁵ Accounts of warrant which include a sensitivity condition have been defended by Fred Dretske, "Conclusive Reasons," *Australasian Journal of Philosophy* 49, no. 1 (1971) and Robert Nozick, *Philosophical Explanations* (Cambridge, MA: Harvard University Press, 1981), 172-79.

⁶ Accounts of warrant which include a safety condition have been defended by Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000) and Duncan Pritchard, *Epistemic Luck* (Oxford: Oxford University Press, 2005); "Anti-Luck Virtue Epistemology," *Journal of Philosophy* 109, no. 3 (2012).

⁷ That sensitivity accounts do not allow for knowledge that we are not brains in vats whereas safety accounts do allow for such knowledge is often held by safety theorists to be a decisive advantage. See Ernest Sosa, "How to Defeat Opposition to Moore,"

Modal reliability conditions for warrant are often formulated in a way that obscures the fact that they face the generality problem. The following is typical:

(MR) S's belief *p* is warranted only if S wouldn't have falsely believed *p* in the nearest world in which *p* is false (SENSITIVITY) / in nearby possible worlds (SAFETY).

To see modal reliabilism in action, consider the following case. Suppose that Lizzy forms the belief that *it is raining outside* on the basis of the testimony of her trustworthy flatmate Laura who has just got in from the pouring rain. Is Lizzie's belief sensitive? Presumably it is, because in the nearest possible world in which it isn't raining, her flatmate Laura doesn't tell her that it is raining and Lizzie doesn't form the belief that it is. Is Lizzie's belief safe? Again, presumably it is, because Laura couldn't easily have been wrong about the weather she recently witnessed and hence there are no nearby worlds in which Lizzie errs in regard to whether it is raining. Hence, Lizzie's belief is safe. So far so good. No mention of belief-forming process types.

However, there are various cases which require (MR) to be finessed in certain ways in order to deliver intuitively correct verdicts. Suppose that Lizzie very nearly consulted her phone's weather app instead of asking her flatmate Laura, and suppose that unbeknownst to Lizzie this weather app has a glitch which causes it to display the forecast for this time last week, when it rained constantly. In nearby worlds in which it isn't raining and in which Lizzie consults the weather app instead of her flatmate, Lizzie winds up believing mistakenly that it is raining outside. What's more, in the nearest world in which it isn't raining, Lizzie's flatmate stayed out longer, and so Lizzie consults her weather app instead of her flatmate. So, in the nearest world in which it is not raining, Lizzie mistakenly believes that it is raining; and in some nearby possible worlds in which it is not raining, Lizzie mistakenly believes that it is raining. Hence, Lizzie's belief is neither sensitive nor safe. But it is counterintuitive to deny Lizzie's belief the status of knowledge. After all, in the actual world her belief is based upon the testimony of a trustworthy informant who told her the truth. Evidently, then, (MR) needs to be refined with reference to the *way* or *method* by which an agent formed her belief:⁸

(MR*) S's belief *p* is warranted only if S wouldn't have falsely believed *p* in the nearest world in which *p* is false (SENSITIVITY) / in nearby possible worlds (SAFETY), *given the way S actually formed her belief*.

What this refinement says is that when we survey the nearby possible worlds/the nearest not-*p* world, although we allow a range of conditions to vary across worlds including the weather, we are to hold fixed the method Lizzie actually used in forming a belief about whether it is raining. But *what* is the method Lizzie used? Precisely *which* features of the situation are we to hold fixed? It is

Philosophical Perspectives 33, no. 13 (1999) and Duncan Pritchard, "Sensitivity, Safety, and Antiluck Epistemology," in *The Oxford Handbook of Skepticism*, ed. John Greco (Oxford: Oxford University Press, 2009).

⁸ As acknowledged by Pritchard, "Anti-Luck Virtue Epistemology," 257-58.

at this point that the generality problem presents itself, because there are different levels of generality at which we can characterize the method of belief-formation that gets held fixed across possible worlds. What we are holding fixed, of course, is the instantiation of a certain type of causal belief-forming process. At the broad end of the spectrum, we could characterize Lizzie's method as *trusting testimony*, and at the narrow end, as something much more specific such as *trusting Laura's testimony about weather conditions which Laura has witnessed in the very recent past*. In between these two extremes are indefinitely many other possible characterizations.

The crucial point is that *which* of these methods (i.e. process types) we select as the relevant one can make a crucial difference to whether the token belief satisfies modal reliability conditions. Suppose we select *trusting testimony* as the relevant method which is to be kept fixed across possible worlds for the purposes of evaluating modal reliability. Accordingly, nearby possible worlds in which Lizzie receives testimony about the weather from other testifiers besides Laura are also counted as relevant for the safety of Lizzie's belief. That significantly increases the scope for failure in regard to the safety condition. But if we select *trusting Laura's testimony about weather conditions which Laura has witnessed in the very recent past* as the method of belief-formation, then that considerably limits the scope for failure in regard to the safety condition, since it takes into account only those worlds in which Lizzie forms a belief by way of one particular person's testimony under quite specific conditions.

Proposals for solving the generality problem

Earl Conee and Richard Feldman have suggested that a satisfactory solution to the generality problem must meet three desiderata.⁹ Firstly, it must be principled. In other words, it must involve the application of a general rule for selecting the relevant process type in any given case; it cannot involve *ad hoc* classifications made on a case-by-case basis. Secondly, it must yield intuitively correct judgments about the level of epistemic warrant enjoyed by any given token belief. Thirdly, it must "remain true to the spirit of the reliabilist approach."¹⁰ That is, it shouldn't rely on non-reliabilist notions such as the strength of an agent's evidence, except insofar as such notions are ultimately spelled out in terms of the success rate of the relevant process type at producing true beliefs across a suitably-defined range of circumstances.

Several proposals have been put forward as to how solve the generality problem. Each of these proposals aims in effect to satisfy the requirements outlined by Conee and Feldman. One could roughly divide these proposals into three camps. Taking their cue from the cognitive sciences, some reliabilists propose that the solution to the generality problem is to be found in the descriptions of belief-forming processes afforded by our best empirical accounts of human

⁹ Earl Conee and Richard Feldman, "The Generality Problem for Reliabilism," *Philosophical Studies* 89 (1998); "Typing Problems," *Philosophy and Phenomenological Research* 65, no. 1 (2002); *Evidentialism: Essays in Epistemology* (Oxford: Oxford University Press, 2004), 135-59.

¹⁰ Conee and Feldman, "The Generality Problem for Reliabilism," 4.

cognition. William Alston and Ralph Baergen have developed this approach.¹¹ The thought is that for any given token belief, empirical accounts of cognition will enable us to fix upon the one relevant process type for evaluating the belief's epistemic status. Another approach, hinted at by Alvin Goldman but more fully developed by Kelly Becker,¹² proposes that the relevant process type is "the narrowest, content-neutral process that is causally operative in belief production."¹³ The idea is that for any given token belief one should select a type that is as narrow as possible within the following two constraints. Firstly, the description of the type must not mention the content of the token belief. So, for Tom's belief that *the sun is shining*, we should not select the type *forming a belief about whether the sun is shining*. Secondly, the description of the relevant type must only make reference to factors that were actually causally relevant to the production of the token belief. So, for instance, the fact that Tom formed his belief on a Tuesday is not causally relevant (let us suppose) and so that fact should be omitted from the description of the relevant process type. Yet another solution that has been put forward in the recent literature is a contextualist one, according to which features of the conversational context determine the truth conditions for a statement of the form "S's belief that *p* was formed by a reliable process." John Greco and Mark Heller have defended this approach,¹⁴ according to which conversational context fixes both the width of the relevant process type and the range of circumstances across which the type success ratio is measured in order for a given belief to count as being produced by a reliable process.

Each of these proposals purports to offer a general rule for identifying the relevant process type for any given token belief. A number of critiques of these proposals are already extant, and I won't repeat them.¹⁵ In essence the objection is that these various proposals fail to do enough to narrow down the range of process types to just one single relevant type for each given token belief. I am sympathetic to this objection. But what I shall argue now is that there are certain sorts of cases which are such that *no* process type that we might single out as the relevant one will deliver an intuitively correct verdict on the epistemic status of the belief.

Complexly-based beliefs and modal reliabilism

By speaking of "complexly-based belief," what I mean to denote is a case in which an agent's belief is based upon inference from a complex evidence set. For the purposes of my argument, I don't need a watertight definition or analysis of complexly-based belief, because I am simply using the term to gesture loosely at a cluster of cases which exhibit features that I shall argue are problematic

¹¹ William P. Alston, "How to Think About Reliability," *Philosophical Topics* 23, no. 1 (1995); Ralph Baergen, *Contemporary Epistemology* (Fort Worth, TX: Harcourt Brace, 1995).

¹² Alvin Goldman, *Epistemology and Cognition* (Cambridge, MA: Harvard University Press, 1986), 50f; Kelly Becker, "Epistemic Luck and the Generality Problem," *Philosophical Studies* 139 (2008).

¹³ *Ibid.*, 363.

¹⁴ Mark Heller, "The Simple Solution to the Problem of Generality," *Noûs* 29, no. 4 (1995); Greco, *Achieving Knowledge*, 76-80.

¹⁵ See especially Conee and Feldman "The Generality Problem for Reliabilism"; "Typing Problems"; but also Richard Swinburne, *Epistemic Justification* (Oxford: Clarendon Press, 2001), 13-20.

for reliabilism. Complexly-based belief is to be contrasted with simply-based belief, such as a belief that *it is 12pm*, based on the evidence of one's wristwatch reading "12pm."

Consider the following fictional case of complexly-based belief. A historian, let's call him Richard, is investigating the question of whether Alberich, a king who at one time reigned over an ancient kingdom called Theusia, was responsible for setting free a people group known as the Mythrians who had once been the slaves of the Theusians. Richard has assembled the following relevant evidence:

- E1: A careful Theusian chronicler writing about a century after the time of Alberich reports that "around a century ago, the Mythrian slaves were freed and thereupon migrated into the region of Dorn, where they made their lasting home."
- E2: Modern archaeological surveys of the area corresponding to the ancient region of Dorn show a sudden appearance of coins and pottery in the archaeological stratum corresponding to the decades following Alberich's reign. Many of the coins and pottery items bear an emblem depicting a king, accompanied by the words "we thank Alberich."
- E3: A fragment of a report from another Theusian chronicler, writing shortly after the reign of Alberich, states that "Alberich will not be remembered kindly by the people of Theusia, for he needlessly gave away the kingdom's most valuable commodities."
- E4: A pamphlet from around the time of Alberich rehearses what appears to be a Theusian pro-slavery slogan: "Mythrian bodies are our commodities."

Richard reflects on this data and comes to the conclusion that Alberich did indeed emancipate the Mythrian slaves. Specifically, he reasons in the following way. To start with, the first chronicler's report, E1, should be regarded as probably accurate (in the absence of contrary evidence) in view of this chronicler's high level of accuracy on matters where historians have been able to independently check his claims. The time period at which E1 says the Mythrians were freed is chronologically consistent with Alberich having freed them. What's more, the report states that the Mythrians migrated into the regions of Dorn and settled there after their emancipation. E2, the sudden appearance of pottery and coins in the region of Dorn in the archaeological stratum corresponding to the period shortly after Alberich's reign, is very probably due to that migration of Mythrians which E1 reports. The fact that many of these coins and pottery items bear emblems expressing gratitude to Alberich strongly indicates that the Mythrians were grateful to Alberich for something of deep importance to them—something important enough for them to have committed it to communal memory by imprinting many of their household items. What might Alberich have done for the Mythrians to make them so enduringly grateful to him? The fact that E1 tells us that the Mythrians were set free from slavery, combined with the fact that only an act so radical as emancipation would be likely to leave a people group so enduringly grateful to the king of the nation that formerly held them captive, makes it pretty probable that Alberich's emancipation of them was the cause of their deep gratitude to him. So, already, just given E1 and

E2, it is probable that Alberich emancipated the Mythrians. E3 and E4, taken together, only serve to increase that probability. E3 states that Alberich gave away the kingdom's most valuable commodities. What commodities might those be? Slaves are typically thought of by their slaveholders as commodities. But this conjecture is confirmed by E4, which shows that pro-slavery writers in Theusia around the time of Alberich specifically used the language of "commodities" to refer to Mythrian slaves, making it quite probable that the "commodities" which E3 blames Alberich for giving away are indeed Mythrian slaves. All in all, then, the cumulative weight of E1-E4 firmly supports the hypothesis that Alberich freed the Mythrian slaves. Richard's conclusion appears strongly warranted. If we suppose that Richard's conclusion is correct, then it is plausible that Richard knows that Alberich freed the slaves.

Now, suppose that if Alberich had *not* freed the Mythrian slaves, his failure to do so would have been due to his being assassinated by a pro-slavery activist before he had the chance to enact his lifelong goal of emancipating the Mythrians. And suppose that if he had been so assassinated, his family would have inscribed his tomb with the words, "Alberich—king, husband, father. His life's purpose was to free the Mythrian slaves." And suppose that if Alberich had been murdered by a pro-slavery assassin, this would have so enraged the Mythrian slaves that they would have revolted and many of them would have successfully escaped Theusia and formed a new community in the region of Dorn, and would have erected stone monuments there declaring "Freedom for Mythria—at last!" A later Theusian chronicler would have reported that the Mythrian slaves gained freedom at around this time. The nearest world to the actual world in which Alberich did *not* free the Mythrian slaves is a world in which the foregoing things occur. Let's call this world *w*. Not only is *w* the nearest world to the actual world in which Alberich didn't free the Mythrians, it is also extremely close to the actual world: the assassination very nearly occurred. In the actual world, the assassin's crossbow bolt missed Alberich's head by mere millimeters. Thus, in a wide range of nearby worlds, Alberich was assassinated and events unfolded as they did in *w*.

In *w* and other nearby worlds in which the assassination attempt succeeded, the historian Richard is confronted with the following pieces of evidence:

- E5: An inscription on the tomb of Alberich, bearing the words: "Alberich—king, husband, father. His life's purpose was to free the Mythrians."
- E6: Several stone monuments in the region of Dorn, dating to shortly after the death of Alberich, inscribed with the words: "Freedom for Mythria—at last!"
- E7: A report from a careful Theusian chronicler writing about a century after the reign of Alberich, stating that "around a century ago, after they gained their freedom, the Mythrians migrated into the region of Dorn."

On the basis of surveying this evidence, Richard concludes that Alberich freed the Mythrian slaves, though he holds his conclusion with slightly less confidence than in the actual world. He reasons that E7 makes the time at which the Mythrian slaves gained freedom chronologically consistent

with Alberich's having freed them, and that this is supported by the existence of E6, the freedom monuments built shortly after Alberich in the region of Dorn, the very region where E7 says some of the Mythrians settled after gaining freedom. The fact that the Mythrians gained freedom around the time of Alberich is well-explained by the hypothesis that Alberich freed them, and this hypothesis is supported by the tomb inscription, E5. The evidence that Alberich freed the slaves is moderately good in world w , but in fact in world w Alberich did not free them.

This scenario creates trouble for both sensitivity and safety accounts of warrant. For sake of ease, let p be the proposition that *Alberich freed the Mythrian slaves*. The historian Richard's actual-world belief in p is counted sensitive provided that Richard refrains from believing p in the nearest world in which p is false and in which Richard uses the same belief-forming method as he did in the actual world. The key question, then, is whether Richard's belief-forming method in world w is the same as his method in the actual world. The answer to this question, of course, hangs on how we characterize the method that Richard uses. And what I want to suggest is that none of the ways of characterizing Richard's method that we might try can give the intuitively correct result that Richard's actual-world belief is an instance of knowledge. The difficulty can be boiled down to the following dilemma:

Either: we characterize Richard's method in a way that doesn't make reference to particular items of evidence, viz., E1-E4. The avoidance of reference to particular items of evidence can be thought of as establishing a minimum level of generality at which we may characterize Richard's method. The narrowest (i.e. least general) method we could select, on this approach, would be something like *inference to the best explanation from archaeological evidence in the region of Dorn from close to the time of the events in conjunction with Theusian manuscript evidence from close to the time of the events*. Given that we avoid reference to particular items of archaeological and manuscript evidence in characterizing Richard's method, we are forced to regard w as a world in which Richard uses the same method as he uses in the actual world. After all, Richard's inference to p from items E5-E7 certainly involves making inferences to the best explanation from both archaeological evidence from Dorn and Theusian manuscript evidence from close to the time of Alberich. Hence, we are forced to regard w as the nearest not- p world in which Richard uses the same method as he uses in the actual world. We are also forced to regard w as a nearby world in which Richard uses the same method as he uses in the actual world. Given that in w Richard's belief p is false, his actual-world belief p fails to meet either the sensitivity condition or the safety condition. Hence, when we characterize the method in terms that fail to mention particular items of evidence, Richard's actual-world belief p is implausibly deemed not to be an instance of knowledge by the sensitivity and safety accounts.

Or: we characterize Richard's method in a way that *does* make reference to particular items of evidence. The inclusion of particular items of evidence in the method can be seen as establishing a maximum level of generality at which we may characterize Richard's method. Given that we include one of the particular items of evidence in our characterization of Richard's method, there is no good reason not to include all of the particular items of evidence that Richard actually used in drawing his conclusion. On this approach, then, we classify Richard's method as *drawing*

inferences from evidence E1-E4, or as something even narrower which makes reference to the particular manner in which Richard evaluated items E1-E4. Given this way of characterizing Richard's method, *w* is *not* counted as a world in which Richard uses the same method as in the actual world, because in world *w* Richard's particular evidence was E5-E7 rather than E1-E4. So, for the purposes of evaluating sensitivity, we are to consider what Richard ends up believing in the nearest not-*p* world in which he possesses items of evidence E1-E4 and draws inferences from these very pieces of evidence. What we are asking, then, is whether Richard would still believe *p* in the nearest world in which he has the very same (excellent) evidence for *p* as he actually has, and yet *p* is false. Let's call this world *x*. The answer to our question is surely that Richard *does* still believe that *p* in such a world. Similarly, if we asked whether I would still believe that I have hands in the nearest world in which I have the same (excellent) visual evidence for that proposition and yet it is false that I have hands, the answer will of course be that I do still believe that I have hands in that world. So, given that we classify Richard's method in a way that keeps fixed the particular (very strong) evidence he actually used, we are forced to conclude that Richard's actual-world belief *p* fails to meet the sensitivity condition and hence, contrary to intuition, that it is not an instance of knowledge.

How about the safety condition? Well, safety only deems relevant those *nearby* worlds in which the agent uses the same method as she actually used. Plausibly, world *x* is not a nearby world. That is, things would have to have gone pretty differently than they actually did in order for a scenario in which Alberich didn't free the slaves to leave behind such a highly misleading set of clues that point so strongly to the conclusion that he did free them: the chronicler's report about the period of Alberich's reign which states that the Mythrian slaves migrated into Dorn after being freed, in conjunction with the pottery and coins appearing shortly after the reign of Alberich in the regions of Dorn with an emblem of a king and the words "we thank Alberich", together with the other chronicler's report which castigates Alberich for giving away the kingdom's most valuable commodities and the pamphlet indicating that the term "commodities" was a way of talking about slaves. Indeed, it is hard to imagine quite how such evidence could have arisen given that Alberich *didn't* free the slaves. That world *x* is distant from the actual world is good news for the safety theorist: it can be ignored for the purposes of evaluating the safety of Richard's actual-world belief that *p*. But so too can almost all of the nearby worlds for the purposes of evaluating safety. The reason is that in almost all of the nearby worlds, the assassination attempt on Alberich succeeded, which means that in all such worlds Richard's evidence base comprises different particulars (viz., E5-E7) than in the actual world. On an approach to method-classification which holds fixed particular items of evidence, Richard's method in the actual world is counted as a different method than the one he uses in all of these nearby worlds in which the assassination attempt succeeded. And so Richard's actual-world belief satisfies the safety condition in virtue of the fact that only *extremely* close worlds are taken to be relevant for determining the safety of his belief, namely, those worlds in which the assassination attempt failed and thus his evidence base contains the very same particulars as in the actual world. In other words, the range of counterfactual circumstances across which we are testing the modal reliability of Richard's method is exceptionally narrow.

It is easy to see why the safety condition is trivialized given an approach to method-classification which holds fixed the detailed particulars of the agent's evidence base. Suppose, for instance, that Susan believes that *Shakespeare = Francis Bacon*, not because she has done any historical research into the matter but because she accepted the say-so of her friend Carmen, who revels in all manner of conspiracy theories. Suppose that in fact it is true that *Shakespeare = Francis Bacon*. It is intuitive that Susan doesn't know this to be so. But given that we hold fixed a method which includes the detailed particulars of Carmen's evidence, namely, *trusting Carmen's testimony that Shakespeare = Francis Bacon* (or something even narrower), there is no nearby world—indeed no possible world at all—in which Susan arrives at a false belief by using this method. Susan's belief is counted as safe on this approach to method-classification which keeps fixed the detailed particulars of the agent's evidence base.

Cases of complexly-based belief like the one involving Richard's historical inferences pose a serious difficulty for modal reliabilism, given that modal reliabilism crucially relies on the notion of a method of belief-formation. What the difficulty boils down to is that in a case of complexly-based belief the intricacies of the agent's evidence base need to be taken into account in the characterization of her belief-forming method otherwise we end up classifying together cases that are really quite different from one another as regards the epistemic status of the agent's belief. But once we take the step of including the intricacies of the agent's evidence base in the characterization of her method, we run into the difficulty that holding fixed such a tightly specified method trivializes the safety condition by drastically narrowing the range of possible worlds which are deemed relevant for evaluating the safety of the agent's belief. As regards sensitivity, holding fixed a highly detailed method of belief-formation will invariably yield the result that the agent's belief is not sensitive, because however good the agent's evidence, she would of course still believe that *p* in the nearest world in which *p* is false and yet she possesses just the same evidence as she does in the actual world.

Complexly-based beliefs and process reliabilism

Let us now turn to consider the situation as regards process reliabilism. Recall that whereas modal reliability is a matter of whether the agent would have avoided error under certain counterfactual circumstances, process reliability is a matter of the success-to-error ratio of the relevant belief-forming process type which is instantiated by the agent's belief, or in other words, a matter of how often the relevant process type yields true beliefs rather than false beliefs across a suitably-defined range of actual and possible circumstances. Of course, there is the far-from-trivial issue of *how* wide a range of circumstances we are to take into consideration when evaluating the truth-ratio of a given process type. For example, if the process type under consideration is *eyesight of medium-sized objects*, it will of course make a big difference to our estimation of the truth-ratio of this process type whether we measure it across a range of circumstances which includes situations where the lighting is poor and the objects are very distant from the eyes or whether instead we measure its

truth-ratio across a narrower range of circumstances in which the lighting is always good and the objects are always close to the eyes. As has been noted elsewhere,¹⁶ this problem of fixing the relevant range of circumstances across which the process type's truth-ratio is measured can be seen as simply another facet of the generality problem. After all, the range of conditions across which a given type's reliability is measured can simply be built into the description of the type itself. For example, rather than simply talking about the reliability of the process type *eyesight*, we can talk about the reliability of the process type *eyesight under such-and-such conditions*. The issue of fixing the appropriate breadth of possible circumstances can simply be subsumed, then, within the larger issue of fixing the appropriate level of generality for evaluating the reliability of a belief-forming process.

Feldman has characterized the generality problem in terms of finding a way of identifying the relevant process type such that the type that is selected is neither too narrow that it is unrepeatable and hence it becomes meaningless to speak of its reliability (the "single case" problem) nor too broad that various token beliefs which vary in their degree of epistemic warrant are all classified under the same relevant process type (the "no distinction" problem).¹⁷ The approach that is implicit in much of the literature which develops reliabilist accounts of warrant is one in which the selection of the relevant process type is guided by intuition on a case-by-case basis. For cases of simply-based belief this approach often works, at least, in the sense that it is typically quite easy in a given case of simply-based belief to identify a process type which is not too narrow as to be unrepeatable and not too broad as to run into the "no distinction" problem and yields an intuitively correct verdict on the level of warrant enjoyed by the agent's belief. Suppose, for instance, that Lawrence arrives at the belief that there are fairies because he really likes the idea and is apt to believe that his fantasies are reality. For this sort of case it is quite easy to fix upon a process type—something like *wishful thinking*—which is not so narrow as to be unrepeatable and yet not so broad as to subsume cases which vary widely in their level of epistemic warrant. The process type *wishful thinking* presumably has a very low truth-ratio, which fits our intuitive judgment that Lawrence's belief has a very low degree of epistemic warrant. Conee and Feldman have argued that this sort of case-by-case approach to selecting the relevant process type is unacceptable in virtue of being unprincipled and ad hoc.¹⁸ However, what I want to suggest is that even if we set aside that rather serious objection, process reliabilism runs into another difficulty when it comes to cases of complexly-based belief such as the one I introduced earlier in the paper: namely, that *no* level of generality that we might select is able to yield intuitively correct verdicts on such cases. Or to put it another way, when it comes to these sorts of cases, there is no space to pass between the twin perils of the "single case" problem and the "no distinction problem."

So let us return to the case of Richard the historian evaluating archaeological and manuscript evidence and arriving at the conclusion that King Alberich liberated the Mythrian slaves. Again, we can draw the following dividing line down the middle of the various process types

¹⁶ See Klemens Kappel, "A Diagnosis and Resolution to the Generality Problem," *Philosophical Studies* 127 (2006).

¹⁷ Richard Feldman, "Reliability and Justification," *The Monist* 68 (1985).

¹⁸ Conee and Feldman, "The Generality Problem for Reliabilism."

instantiated by Richard's belief: *either* we select a type which makes no mention of the particular items of evidence on which Richard based his belief, *or* we select a type which does mention the particulars. Let's begin with the first approach. Again, the narrowest we can go without mentioning evidential particulars is something like the type *inference to the best explanation from archaeological evidence in the region of Dorn from close to the time of the events in conjunction with Theusian manuscript evidence from close to the time of the events*. The problem is that this type is too broad: it is instantiated by token beliefs which vary in their level of epistemic warrant. This process is instantiated by Richard's actual belief that *p*, which is based upon excellent evidence (namely, E1-E4). But it is also instantiated by Richard's belief that *p* in the many nearby worlds in which Alberich was assassinated before he could free the slaves and in which the slaves rebelled and then migrated into the region of Dorn. In these worlds, Richard arrives at the belief that *p* on the basis of moderately good evidence (namely, E5-E7). Intuitively, there is a difference between the level of warrant enjoyed by Richard's actual-world belief, based as it is on excellent evidence, and the level of warrant enjoyed by his belief in a range of nearby worlds, which is based on moderately good evidence. And yet, we shall be forced to say that these beliefs all have the very same degree of warrant as one another if our selection of the relevant process type involves no mention of the particular items of evidence which were used.

On the other hand, if we opt to select a process type which makes mention of evidential particulars, the broadest type we can select is something like *drawing inferences from E1-E4*. Naturally we avoid the problem we encountered with the first approach, given that Richard's actual-world belief instantiates this process type but his belief in nearby worlds which is based on E5-E7 does *not* instantiate this type. The problem we now encounter is that the process type we have selected is too narrow; it is doubtful that we can meaningfully speak of its reliability. We run into the single case problem. As Feldman articulates it, the single case problem is the problem that "if relevant types are characterized very narrowly then the relevant type for some or all process tokens will have only one instance (namely, that token itself)."¹⁹ But Kelly Becker rightly points out that this way of putting things leaves it unclear whether this is supposed to be "a problem for any process actually used only once, no matter how individuated, or...a problem just for those processes so narrowly individuated that they *could* be used only once."²⁰ Becker is inclined to think that the latter is the real source of trouble. Now, in the way I handled things on the former approach (the approach on which we select a process type that makes no mention of evidential particulars), I already assumed that a process type could be considered repeatable even if it is used only once in the actual world, provided that it could in principle be used in other possible circumstances than the precise ones in which it is actually used.²¹ The process type *drawing inferences from E1-E4* is only used once in the actual world, we may suppose. But it can be considered repeatable provided there are possible worlds in which this process type is employed under circumstances which differ somewhat from those in which it is actually used. The difficulty,

¹⁹ Feldman, "Reliability and Justification," 160-61.

²⁰ Becker, "Epistemic Luck and the Generality Problem," 361.

²¹ Becker (Ibid.) argues for this understanding of repeatability.

however, is that because Richard's having the very particular evidence base he actually has (viz., E1-E4) is so exquisitely sensitive to the way that events unfolded around the reign of Alberich, we simply can't vary the circumstances very much at all without the result that Richard winds up with somewhat different evidence. In other words, the range of possible circumstances in which Richard counts as using the type *drawing inferences from E1-E4* is extremely narrow; too narrow for it to be meaningful to speak of the type's being either "reliable" or "unreliable" across such a narrow range of possible circumstances.

When we talk about a belief-forming process being "reliable," we mean that the process type in question achieves a high success rate when deployed across a certain range of circumstances. The notion of reliability is trivialized if the range of circumstances across which we measure the success rate of a process type is extremely small. What counts as too small is surely vague, but I think we have an intuitive sense that it ceases to be meaningful to speak of "reliability" if the range of circumstances across which we measure the type's success rate is extremely narrow. This point applies not just to the reliability of belief-forming processes, but also to reliability discourse more generally. Would we wish to come to a judgment about whether a certain car is reliable at starting if we are only allowed to take into account its success rate at starting when it is inside a garage where the temperature is between 18 and 22 degrees Celsius and when the car has just been given a full checkup by a mechanic? Probably not.

More generally, the situation will be like this for many cases of complexly-based belief. The reason for this, in short, is that these sorts of cases involve an agent basing her belief on a detailed evidence base, which is such that the details *really matter* when it comes to evaluating the belief's degree of warrant. What's more, in these sorts of cases, very slight changes in the circumstances that yielded the agent's very particular evidence base will lead to the agent's having a somewhat different set of detailed evidence. This creates the following dilemma for the process reliabilist. Either the process reliabilist characterizes the agent's belief-forming process in a way that glosses over the particular details of the agent's evidence base, which leads to the result that cases that intuitively differ as regards their degree of warrant end up being classified under the same relevant process type, yielding the counterintuitive judgment that these various beliefs possess the same degree of warrant. Or, the process reliabilist characterizes the agent's belief-forming process in a way that takes account of the very particular details of the agent's evidence base, but then runs into the problem that the range of circumstances in which such a detailed process type can be employed is exceptionally narrow; too narrow for it to be meaningful to speak of the type's reliability. In short, when it comes to these sorts of cases, there is no level of generality to be found which avoids both the no distinction problem and the single case problem.

Conclusion

I have argued that certain cases involving what I have termed "complexly-based belief" pose a serious difficulty for both process reliabilism and modal reliabilism. What characterizes these cases

is that they involve complex inferences to the best explanation. The vast majority of our scientific, philosophical, and historical beliefs are of this sort. The problem is that with such beliefs as these we cannot identify *any* process type which is at the right level of generality so as to yield an intuitively correct verdict on the epistemic status of the belief. As we saw earlier, various proposals have been put forward in the literature for solving the generality problem. Objections to these solutions typically allege that these solutions don't really narrow things down to just one process type for any given token belief. But if the thesis of this paper is correct, then the situation for reliabilism is considerably worse than that, because any solution to the generality problem will be premised upon the assumption that for any given token belief there is *at least one* process type which, if selected as the relevant type, would yield an intuitively correct verdict on the epistemic status of the token belief. If that is not in fact the case, then the generality problem is insoluble. The reliabilist may then retreat to the claim that reliabilism gives an adequate account of epistemic warrant only for cases of simply-based belief—cases of forming a belief on the basis of the checking of a clock, or a glance out of the window, and other similar examples which are the staple of reliabilist epistemologies—but concede that, as Nicholas Wolterstorff puts it, “[T]he very idea of identifying methods of belief-formation and determining their reliability has little application when it comes to the deep questions of philosophy and religion, and of many other areas of human life. The idea lacks purchase.”²²

²² Nicholas Wolterstorff, “The Significance of Inexplicable Disagreement,” in *Religious Faith and Intellectual Virtue*, ed. Laura Frances Callahan and Timothy O'Connor (Oxford: Oxford University Press, 2014), 327.