

Process Reliabilism's Troubles with Defeat

Bob Beddor

Published in *The Philosophical Quarterly*, 2015

Abstract

One attractive feature of process reliabilism is its reductive potential: it promises to explain justification in entirely non-epistemic terms. In this paper, I argue that the phenomenon of epistemic defeat poses a serious challenge for process reliabilism's reductive ambitions. The standard process reliabilist analysis of defeat is the "Alternative Reliable Process Account" (ARP). According to ARP, whether S's belief is defeated depends on whether S has certain reliable processes available to her which, if they had been used, would have resulted in S not holding the belief in question. Unfortunately, ARP proves untenable. I show, by way of counterexample, that ARP fails to articulate either necessary or sufficient conditions on defeat. Process reliabilists must either provide an alternative reductive account of defeat or renounce their reductive aspirations.

1. Introduction

Since the publication of Goldman's "What is Justified Belief?" in 1979, process reliabilism has been one of the main accounts of the nature of epistemic justification, and it remains an active research programme. One of the primary attractions of process reliabilism is that it purports to provide a reductive account of epistemic justification: it purports to spell out what it is for a belief to be justified without using any epistemic notions.¹

In this paper, I discuss a challenge for process reliabilist accounts of justification – one that has received insufficient attention in the literature. On pain of counterexample, process reliabilists need to invoke a "No Defeaters" clause. What's more, in order to fulfil their reductive ambitions, they need to give an account of defeat that's free from any epistemic terms.

To their credit, process reliabilists have tried to provide such an account. The standard process reliabilist line spells out defeat in counterfactual terms: it says that whether S's

¹The opening pages of Goldman (1979) contain a clear statement of this reductive ambition: "I want a theory of justified belief to specify in non-epistemic terms when a belief is justified." (1979: 90) In the subsequent literature, many have taken this ambition to be central to the spirit of reliabilism (see e.g. Conee and Feldman 1998: 4-5).

Why want a reductive account of justification in the first place? One potential reason is that reductive accounts can explain why justificatory properties supervene on non-epistemic properties: if all it is for a belief to be justified is for it to satisfy non-epistemic conditions C_1 - C_n , then any two beliefs that have the same non-epistemic properties will have the same justificatory status.

belief that p is defeated depends on whether there are any reliable processes available to S which, if they had been used, would have resulted in S 's not believing p . Following Lyons (2009, forthcoming), I'll call this the "Alternative Reliable Process Account" (ARP). In this paper, I argue that ARP is subject to counterexample. It remains to be seen whether process reliabilists can offer some alternative reductive account of defeat that fares better.²

This paper is structured as follows. In §2, I explain why process reliabilists need an account of defeat. In §3, I introduce ARP and illustrate how it handles certain cases of defeat. The rest of the paper is devoted to causing trouble for ARP: §4 offers an example that shows that ARP fails to articulate sufficient conditions for defeat; §5 offers an example that shows that ARP fails to articulate necessary conditions for defeat. I conclude that process reliabilists must jettison ARP in favour of some other account of defeat.

2. Why Process Reliabilists Need a Story About Defeat

To see why process reliabilists need a story about defeat, consider a simple version of process reliabilism that lacks any "No Defeaters" clause:

Simple Process Reliabilism: S 's belief that p is justified at t iff S 's belief that p at t is the result of a reliable belief-forming process (or reliable belief-forming processes).³

To see why such an account proves inadequate, consider a stock example of defeat:

Consuela and the Vase: Consuela sees a red vase in good lighting conditions at t_1 . Consequently, she comes to believe RED: ⟨There's a red vase in front of me⟩. At t_2 , a usually reliable informant tells Consuela that she's actually looking at a white vase illuminated by a red light. Consuela has no reason to distrust this informant; nonetheless, she disregards his testimony and continues to believe RED.

Fill in the details in the right way, and most people have the intuition that Consuela's belief in RED is defeated at t_2 : though it may be *prima facie* justified, it's not *ultima facie* justified (that is, it's not justified *full stop*). But Simple Process Reliabilism predicts precisely the opposite. After all, Consuela's belief in RED at t_2 is the result of vision operating in good lighting conditions – a paradigmatic example of a reliable cognitive process.

²Arguably, process reliabilists aren't the only ones in trouble here. Various approaches to justification require a "No Defeaters" condition. (Consider, for example, a simple version of phenomenal conservatism, according to which S is justified in believing p iff it seems to S that p and S lacks any defeaters for believing p . Cf. Huemer 2007.) In so far as they share process reliabilists' reductive goals, proponents of any such approach will need to analyse defeat in non-epistemic terms.

For the purposes of this paper, I restrict my attention to process reliabilism, since process reliabilists have been particularly explicit in both their reductive ambitions and their proposed treatment of defeat.

³Despite the amount of press it's received, it's not clear whether anyone has endorsed Simple Process Reliabilism, so it may be something of a strawperson. Nonetheless, it provides a useful entry point for exploring process reliabilists' difficulties with defeat.

Is there any way to defend Simple Process Reliabilism from this objection? One might be inclined to reply as follows:

“Everyone knows that process reliabilists face the Generality Problem: they face the notoriously difficult problem of typing belief-forming processes.⁴ You’ve picked a fairly coarse-grained way of typing Consuela’s belief-forming process at t_2 : you’ve characterized it as vision (or vision operating in good lighting conditions). But perhaps we should appeal to a more fine-grained characterisation of the process responsible for Consuela’s belief. Let T be a function that takes as input visual experiences *together with testimony that those experiences are misleading* and produces as output belief in the content of those experiences. If we take T to be the correct way of typing Consuela’s belief-forming process, it’s not clear that her belief in RED was formed by a reliable process.”

Call this the “Typing Defence.” In evaluating the Typing Defence, we should first note that Consuela’s case is just one among many. Indeed, we can concoct an infinite variety of cases with the following structure: *S forms an intuitively justified belief that p at time t_1 . Then, at t_2 , S acquires a good reason to abandon her belief that p. Nonetheless, S sticks to her guns and continues to believe p anyway.* And so the proponent of the Typing Defence has her work cut out for her: for each such case, her proposed method for typing belief-forming processes will need to deliver the result that S’s belief-forming process at t_1 is reliable and that S’s belief-forming process at t_2 is unreliable.

If we turn to some of the leading attempts to solve the Generality Problem in the literature, we find that they don’t deliver these results. For example, one popular approach to the Generality Problem is to type belief-forming processes in accordance with the “common sense” classifications that the folk employ in ordinary life.⁵ Any such constraint on typing belief-forming processes bodes ill for the Typing Defence. After all, the folk are far more likely to classify belief-forming processes as instances of *vision* or *vision operating in good lighting conditions* than they are to classify belief-forming processes as instances of *vision together with testimony that those visual experiences are misleading*.⁶

Of course, one might respond by simply rejecting such approaches: one might insist that the only adequate solution to the Generality Problem is a solution that’s consistent with the Typing Defence. But I see no reason to impose such a stringent adequacy condition on solutions to the Generality Problem. Typing belief-forming processes is difficult enough without worrying about defeat. Thus it would be nice if process reliabilists could give an independent treatment of defeat – a treatment of defeat that’s compatible with a

⁴For a seminal statement of the Generality Problem, see Conee and Feldman (1998).

⁵Some of Goldman’s remarks in his (1979) paper suggest an approach along these lines. Jönsson (2013, forthcoming) and Olsson (forthcoming) both defend versions of this approach.

⁶Another attractive approach to the Generality Problem is to type belief-forming processes causally. (See e.g. Goldman 1986; Becker 2008.) A natural way of developing a causal approach to the Generality Problem is to insist that if a feature f doesn’t causally affect whether S believes p at t , our way of typing the belief-forming process responsible for S’s belief that p at t shouldn’t mention f . Any such solution also stands in tension with the Typing Defence, since there’s no guarantee that defeaters will always causally affect the target belief. For example, in the case of Consuela and the Vase, we can imagine Consuela (unjustifiably) regards her interlocutor as completely unreliable; hence his testimony doesn’t causally affect her credence in RED.

variety of solutions to the Generality Problem.

Most process reliabilists agree with me on this front. Most process reliabilists reject Simple Process Reliabilism, and many of them do so – at least in part – because of Simple Process Reliabilism’s acknowledged difficulties handling cases of defeat.⁷ Instead of Simple Process Reliabilism, most process reliabilists opt for what we can call a “Two Step” version of process reliabilism.

By a “Two Step” theory of justification, I mean any theory that comprises two separate components: a *prima facie* justification condition and a “No Defeaters” condition. Process reliabilists who go for a Two Step theory typically take Simple Process Reliabilism (or some close variant) and convert it into an account of *prima facie* justification. For instance, the following is a straightforward version of Two Step Process Reliabilism:

- S’s belief that *p* is (*ultima facie*) justified at *t* iff
- (i) S’s belief that *p* is the result of a reliable belief-forming process (or belief-forming processes) at *t*
 - (ii) S’s belief that *p* isn’t defeated at *t*.

A Two Step process reliabilist will allow that Consuela’s belief in RED is *prima facie* justified at *t*₂, but deny that it’s *ultima facie* justified at *t*₂, since it doesn’t satisfy condition (ii).⁸

However, process reliabilists who go down this road face an obvious challenge. *Defeat* is an epistemic notion *par excellence*. Thus, on pain of abandoning their reductive aspirations, Two Step process reliabilists need to provide a reductive account of defeat.

3. The Alternative Reliable Process Account of Defeat

To their credit, process reliabilists have acknowledged this obligation and tried to discharge it. The classic process reliabilist story about defeat was first proposed by Goldman (1979), and has found a recent champion in Lyons (2009, forthcoming). It goes like this:

Alternative Reliable Process Account of Defeat (ARP): S’s belief that *p* is defeated at *t* iff there’s some reliable or conditionally reliable process available to S at *t* which, if it had been used by S in addition to the process actually used, would have resulted in S’s not believing *p* at *t*.⁹

⁷See e.g. Goldman (1979); Lyons (2009, forthcoming); Grundmann (2009); Bedke (2010).

⁸In addition to its superior handling of Consuela and the Vase, Two Step Process Reliabilism may fare better than Simple Process Reliabilism when it comes to dealing with Bonjour’s case of Norman the clairvoyant (Bonjour 1985) and Lehrer’s case of Truetemp (Lehrer 1990). As Goldman (1986) suggests, one option is to say that these characters’ beliefs are *prima facie* justified but defeated. (Though see Lyons (2009) for a version of process reliabilism that precludes Norman and Truetemp’s beliefs from possessing *prima facie* justification.)

⁹Grundmann (2009) and Bedke (2010) defend similar counterfactual accounts of defeat. While their accounts (particularly Grundmann’s) differ from ARP in some important details, I’ll forego a discussion of such details since – differences notwithstanding – the counterexamples I offer to ARP straightforwardly generalize to their proposals.

How does ARP handle Consuela and the Vase? I take it that proponents of ARP will offer the following diagnosis: Consuela's belief in RED is defeated at t_2 because she could have given weight to her interlocutor's testimony. More precisely: there's a possible process that takes his testimony as input and spits out a fairly high credence in the content of that testimony. Plausibly, this process is either reliable or conditionally reliable (or at least there's a natural way of typing this process that makes it either reliable or conditionally reliable). Consuela could have used this process in addition to visual perception; had she done so, she would have abandoned her belief in RED.

I'm happy to grant that ARP is capable of explaining our intuition about Consuela and the Vase. Unfortunately, there are at least some cases where ARP clearly delivers the wrong results.

For ease of exposition, it will be useful to explicitly distinguish ARP's sufficiency condition for defeat from ARP's necessity condition:

Sufficiency: If there's some reliable or conditionally reliable process available to S at t which, if it had been used by S in addition to the process actually used, would have resulted in S's not believing p at t , then S's belief that p is defeated at t .

Necessity: If S's belief that p is defeated at t , there's some reliable or conditionally reliable process available to S at t which, if it had been used by S in addition to the process actually used, would have resulted in S's not believing p at t .

In what follows, I first present a counterexample to Sufficiency (§4); next, I offer a counterexample to Necessity (§5).

4. A Counterexample to Sufficiency

Before presenting what I take to be a convincing counterexample to Sufficiency, I'll briefly mention an unconvincing counterexample. Kvanvig (2007) asks us to imagine that "[T]here is a competent cogniser who disagrees with you about something you know to be true... There is a reliable process which if you had used it would have resulted in a different belief: namely, ask this cogniser and believe what is reported." (2007: 1-2)

The reason I take Kvanvig's counterexample to be unconvincing is that Goldman's initial (1979) discussion of ARP makes it clear he doesn't intend the alternative processes in question to include consulting new interlocutors. Goldman writes:

[I]t seems implausible to say all 'available' processes ought to be used, at least if we include such processes as gathering new evidence. Surely a belief can sometimes be justified even if additional evidence-gathering would yield a different doxastic attitude. What I think we should have in mind here are such additional process as calling previously acquired evidence to mind, assessing the implications of that evidence, etc. (Goldman 1979)

It seems clear that Goldman intends ARP to be restricted to alternative reliable cognitive processes that are in an important sense internal; they do not involve further research. And so the process suggested by Kvanvig – *consulting competent cognizers who disagree with you* – won’t qualify as “available” in the relevant sense.¹⁰

Nonetheless, I think there are other cases with a similar structure that do present compelling counterexamples to Sufficiency. Here’s one:

Thinking About Unger: Harry sees a tree in front of him at t . Consequently, he comes to believe the proposition TREE: ⟨There is a tree in front of me⟩ at t . Now, Harry happens to be very good at forming beliefs about what Peter Unger’s 1975 time-slice would advise one to believe in any situation. Call this cognitive process his “Unger Predictor”: for any proposition p , any agent A , and any situation s , Harry’s Unger Predictor spits out an accurate belief about what doxastic attitude Unger’s 1975 time-slice would advise A to take towards p in s .

What’s more, Harry has a very high opinion of Unger’s 1975 time-slice: whenever it occurs to Harry that Unger would advise him (Harry) to suspend judgement about p , this causes Harry to become chagrined and suspend judgement about p . So if Harry had used his Unger Predictor, he would have come to believe SUSPEND: ⟨Unger would advise me (Harry) to suspend judgement regarding TREE⟩. This would, in turn, have caused Harry to suspend judgement regarding TREE.

Sufficiency predicts that Harry’s belief in TREE is defeated at t . After all, his Unger Predictor is a reliable process: it systematically produces true beliefs about the advice of Unger’s 1975 time-slice. It’s also an *internal, cognitive* process that’s available to him at t . So there’s a reliable, internal belief-forming process that’s available to Harry at t , which, if it had been used by Harry, would have resulted in him not believing TREE at t .

However, this verdict about the case strikes me as wrong: my intuition is that Harry’s belief in TREE is justified at t . After all, at t , Harry isn’t thinking about Unger, or entertaining any sceptical doubts; indeed, we can stipulate he hasn’t engaged in sceptical ruminations in a very long time. At t , it seems that he has every reason to think that there’s a tree in front of him and no good reason to suspend judgement. The mere fact

¹⁰Given this restriction of what counts as an alternative reliable process “available” to the agent, it’s not clear that Goldman is entitled to his later verdict on Kornblith’s (1983) case of Jones, “a headstrong young physicist” who gives a talk wherein he announces his belief that p . Jones is unable to withstand any sort of criticism, and consequently doesn’t listen to his colleague’s devastating objection. (It’s crucial to Kornblith’s case that Jones doesn’t even hear the colleague’s objection, and so does not possess the counterevidence that it furnishes.) Goldman (1992) seems to concur with Kornblith’s assessment that the physicist’s belief that p is defeated in this case. But it’s not clear that ARP can deliver this verdict, since the reliable process involved in listening to his colleague’s testimony wouldn’t be entirely internal – it would, in an important sense, involve engaging in further research.

I won’t press this objection, however, since I don’t find Kornblith’s case entirely convincing. While I certainly concur that there’s something epistemically defective about the physicist’s overall state, it’s not clear to me that this defect defeats his belief that p . (It seems plausible to me that a person can have justified beliefs, even though those beliefs are sustained, at least in part, by a variety of epistemic vices.)

that *if* he were to engage in reflection about what Unger would think about his situation, he would suspend judgement regarding TREE shouldn't deprive his belief of justification.

Thus I take it that Thinking About Unger is a counterexample to Sufficiency, and hence ARP. Is there any way of repairing ARP to get around this counterexample? When faced with this case, it's natural to feel that there's something fishy about the connection between Harry's Unger Predictor and his belief in TREE. And thus it's natural to suspect that if we can flesh out this worry – if we can articulate in what sense this connection is fishy – we'll be able to revise ARP in a way that avoids generating the prediction that Harry's belief in TREE is defeated.

So what exactly is fishy about the connection between Harry's Unger Predictor and his belief in TREE? It seems to me that there are two natural hypotheses. The first is that *causal indirectness* is to blame. If Harry were to use his Unger Predictor, this wouldn't directly cause him to suspend judgement about TREE. The only immediate effect of his Unger Predictor would be a belief in SUSPEND; the actual suspension of judgement regarding TREE would occur as a further, downstream effect.

A second, closely related hypothesis is that a subject matter mismatch is the culprit. According to this hypothesis, Harry's Unger Predictor doesn't produce beliefs about the right sort of subject matter: it produces beliefs about what Unger would advise, not about whether there are trees in front of Harry.

These two hypotheses suggest two different ways of trying to repair ARP. The first hypothesis suggests a version of ARP that includes a causal directness requirement:

ARP*: S's belief that p is defeated at t iff there's some reliable or conditionally reliable process available to S at t which, if it had been used by S in addition to the process actually used, would have directly resulted in S's not believing p at t .

The second hypothesis suggests a version of ARP that includes a subject matter requirement. A natural way of spelling out this requirement is to insist that the alternative process must produce a doxastic attitude towards p :

ARP**: S's belief that p is defeated at t iff there's some reliable or conditionally reliable process X available to S at t such that:
(i) The output of X is a doxastic attitude towards p
(ii) If X had been used by S in addition to the process actually used, it would have resulted in S's not believing p at t .

Will either ARP* or ARP** work? I think not. To see why, imagine that Harry had used his Unger Predictor, and hence had come to suspend judgement regarding TREE. Now, presumably there is *some* process that directly results in this suspension of judgement; it's just a two-component process. The first component is his Unger Predictor, which outputs a belief in SUSPEND; the second component is a process we can call his "Unger Implementer": it takes as input a belief in SUSPEND, and produces suspension of judgement

regarding TREE. Call this two-component process Harry’s “Unger Emulator” (depicted in fig. 1).

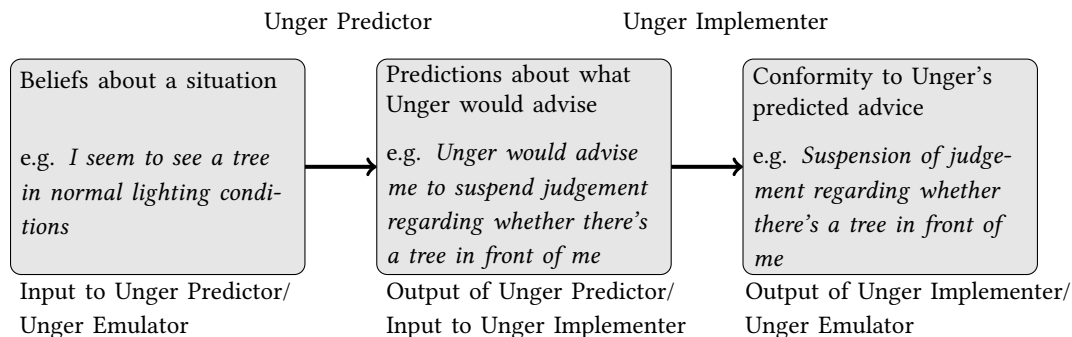


Figure 1: Unger Emulator

Now, Harry’s Unger Emulator is a reliable process. After all, it produces no beliefs; *a fortiori*, it doesn’t produce any false beliefs. So it seems Harry has an alternative reliable belief-forming process available to him (his Unger Emulator) which, if he had used it in addition to the process actually used (visual perception), would have *directly* resulted in him suspending judgement about TREE. Hence, ARP* delivers the result that his belief in TREE is defeated at t .¹¹ What’s more, the output of his Unger Emulator is a doxastic attitude towards TREE (suspension of judgement is, after all, a doxastic attitude); hence ARP** also delivers the verdict that his belief in TREE is defeated at t .

One might object to my claim that Harry’s Unger Emulator is a reliable belief-forming process. There are at least two ways one might try to motivate this objection. First, one could point out that since the Unger Emulator doesn’t produce any beliefs, it doesn’t count as a *belief-forming* process at all. Second, one could dispute the assumption that failure to produce false beliefs is sufficient for reliability. On a sufficiently nuanced conception of reliability, suspending judgement on a truth will count as worse than believing a truth.¹² Since the Unger Emulator suspends judgements on all truths, it’s not a particularly reliable process.

But even if we grant this objection, we can easily amend the case. Imagine an epistemologist – call her Shmunger – whose scepticism is much more circumscribed than that of

¹¹Could one insist that, strictly speaking, the process that would have *directly* resulted in his suspension of judgement wouldn’t have been his Unger Emulator; rather, it would have been just the second stage of the Unger Emulator (the Unger Implementer)? Even if this view could be motivated (which I doubt), it will be of no help to the proponent of ARP*; after all, the Unger Implementer produces no beliefs either, hence it too is reliable.

¹²The best candidates for nuanced conceptions of reliability along these lines come from the literature on scoring rules, where the discussion is usually couched in terms of credences rather than the tripartite distinction between belief, suspension, and disbelief. See, for example, Joyce (1998); Gibbard (2007); Moss (2011). (Goldman also discusses such scoring rules in places (Goldman and Shaked 1991, Goldman 1999, Goldman 2010), though he presents such scoring rules as measures of “degrees of truth possession” or “veritistic value” rather than reliability.)

Unger’s 1975 time-slice. Shmunger has lots of true beliefs about math, physics, history – you name it. Shmunger is only a sceptic about trees: in any given situation, she suspends judgement on whether trees are present.

Suppose, furthermore, that Shmunger wants others to share her doxastic attitudes: she wants others to hold her (entirely accurate) beliefs about math, physics, etc.; she also wants everyone to share her arboreal scepticism. Now we can reproduce the counterexample using Shmunger:

Thinking About Shmunger: Harry’s twin Larry has an extremely reliable Shmunger Predictor: given any scenario, Larry’s Shmunger Predictor reliably predicts what Shmunger would advise Larry to believe in that scenario. Larry’s Shmunger predictor is also part of a Shmunger Emulator: whenever Larry predicts that Shmunger would advise Larry to take up a particular doxastic attitude towards a proposition, he invariably takes up that attitude. Now Larry sees a tree in normal lighting conditions. Since Larry isn’t thinking about Shmunger, he comes to believe TREE: ⟨There is a tree in front of me⟩. But if he had used his Shmunger Emulator, he wouldn’t have done so.

Clearly, Larry’s Shmunger Emulator is a belief-forming process: it produces lots of beliefs on sundry topics. What’s more, the beliefs in question are always true. Consequently, it seems to qualify as a reliable belief-forming process, even given a suitably nuanced conception of reliability. If Larry had used this process, it would have *directly* resulted in Larry’s suspending judgement about TREE. So ARP* predicts that Larry’s belief that there’s a tree in front of him is defeated. And since the output of Larry’s Shmunger Emulator is a doxastic attitude towards TREE, ARP** delivers the same prediction.

So it seems that the most natural ways of revising ARP to get around the counterexamples to Sufficiency prove unsuccessful.¹³ And it’s worth noting that nothing hinges on

¹³For those who suspected that some sort of subject matter mismatch was to blame, a natural next attempt at patching ARP would be the following:

ARP*:** S’s belief that p is defeated at t iff there’s some process X available to S at t such that:
(i) X is reliable about p -related matters
(ii) If X had been used by S in addition to the process actually used, it would have resulted in S’s not believing p at t .

If we coupled ARP*** with a nuanced conception of reliability, we could capture the intuition that Larry’s belief in TREE is justified. After all, though his Shmunger Emulator is generally reliable, it isn’t reliable about TREE-related matters: it systematically recommends suspending judgements about the presence of trees, hence isn’t reliable about their presence.

However, proponents of this account face the difficult task of specifying what counts as a “ p -related matter.” (As an anonymous referee helpfully observed, this could be viewed as a “second round” of the Generality Problem.) If in Thinking About Shmunger we take a “TREE-related matter” to include any proposition about the presence of trees (as the foregoing paragraph suggests), then the envisioned fix won’t be able to handle cases of even more circumscribed scepticism. (Imagine Shmunger isn’t a wholesale sceptic about the presence of trees; indeed, she’s generally quite reliable about the presence of trees. She’s only a sceptic about the presence of, say, birch trees when viewed from such-and-such an angle.) To deal with such cases, we’d need to

the details of Thinking About Unger/Shmunger; we can construct many more counterexamples with the same general structure. Here's the formula:

Step 1: Describe a case where an agent S has a *prima facie* justified belief that *p*.

Step 2: Endow S with a dormant reliable belief-forming process X that meets the following conditions:

- (i) If S were to use X, the output of X would be suspension of judgement about *p*
- (ii) The mere availability of X does not, intuitively, give S a good reason to cease believing *p*.

What's more, even if we can revise ARP in a way that avoids every case of this form, ARP won't be in the clear: as we're about to see, there are also counterexamples to Necessity. Any satisfactory revision of ARP will have to deal with these as well.

5. A Counterexample to Necessity

Here's a counterexample that shows that ARP fails to articulate necessary conditions for defeat:

Job Opening: Masha tells Clarence that her department will have a job opening in the fall. Clarence believes Masha; assuming that Masha is usually reliable, Clarence's belief counts as *prima facie* justified. Sometime later, Clarence speaks with the head of Masha's department, Victor, who informs him that the job search was cancelled due to budget constraints. Now suppose that Clarence harbors a deep-seated hatred of Victor that causes him to disbelieve everything that Victor says; what's more, no amount of rational reflection would rid Clarence of this inveterate distrust. Consequently, he continues to believe that there will be a job opening in the fall.

Intuitively, Clarence's belief that there will be a job opening in the fall is defeated.¹⁴ But it doesn't seem that Necessity delivers this result. Consider: is there some alternative reliable process available to Clarence such that, if Clarence had used it in addition to the process he actually used (namely, trusting Masha's testimony), he wouldn't have believed that there will be a job opening in the fall? It doesn't seem so. After all, we've stipulated that Clarence has an intractable distrust of Victor – one that no amount of reflection or therapy could dislodge. So it doesn't seem that there's any process available to Clarence that takes Victor's testimony as input and produces a fairly high credence in said testimony.

give a more restrictive account of "TREE-related matters." But the more we restrict what counts as a "*p*-related matter", the more difficulty we'll have capturing intuitions about mundane cases of defeat. For instance, in Consuela and the Vase, we'd like to say that Consuela's belief in RED is defeated even though her interlocutor isn't reliable about whether the particular vase Consuela is looking at is red.

¹⁴It would be different if Clarence had good reason to deem Victor's testimony unreliable, but in our scenario this isn't the case.

One might insist that Clarence has a very general process (GP) available to him, which takes any testimony t he's received as input and produces an increased credence in t as output. It's just that Clarence isn't capable of employing this process for every value of t . But even if we grant that he has such a process available to him, this won't be enough to rescue Necessity. In order to rescue Necessity, the following counterfactual would need to be true:

If Clarence had used GP, he wouldn't have believed there will be a job opening in the fall.

But given the way the case is set up, this counterfactual is false. After all, Clarence has only received two pieces of testimony regarding the potential job opening: Masha's and Victor's. He can't plug Victor's testimony into GP, and plugging Masha's testimony into GP was what caused him to believe that there will be a job opening in the first place.

Of course, one could part ways with Goldman and opt for a less restrictive conception of the belief-forming processes "available" to an agent. Someone who takes this approach could point out that Clarence is capable of asking other people whether the department will have a job opening in the fall; if Clarence had pursued such inquiries, he would have presumably acquired independent evidence that the job search was cancelled. Our defender of Necessity could then insist that this suffices for Clarence to have an alternative reliable process available to him which, if used in addition to the process actually used, would have resulted in him abandoning the belief that there will be a job opening in the fall.

However, there are at least two problems with this manoeuvre. First, we can set up the case in such a way that Clarence's independent investigations would have been fruitless: simply stipulate that, despite persistent inquiries, Clarence would have been unable to unearth any further evidence regarding the potential job opening. Second, a process reliabilist who makes this manoeuvre will open herself up to a host of additional counterexamples to Sufficiency. For instance, she'll fall prey to Kvanvig's counterexample (discussed in §4): if engaging in further inquiry counts as a belief-forming process available to an agent, Sufficiency entails that S's belief that p will be defeated whenever there's some reliable cognizer with whom S is acquainted who believes $\neg p$.

Here too nothing hinges on the details of case; it's easy to cook up similar examples. Here's the recipe:

Step 1: Describe a case where an agent S has a *prima facie* justified belief that p .

Step 2: Stipulate that S receives strong evidence e that p is false.

Step 3: Stipulate that, due to hatred, prejudice, or just some psychological quirk, S is cognitively incapable of responding appropriately to e , and hence persists in believing p .

Every case we cook via this recipe will put defenders of Necessity in a similar bind. For every such case, defenders of Necessity may be tempted to insist that even though

S is unable to respond appropriately to e , S is capable of pursuing further inquiries into the truth of p ; had she done so, she would have unearthed some further evidence e^* that would have caused her to suspend judgement on p . This, they may be tempted to insist, is all that it takes for S to satisfy Necessity, given a suitably liberal conception of “availability.” But broadening our conception of availability will invariably lead from frying pan to fire: any such broadened conception of availability will give rise to legions of Kvanvig-style counterexamples to Sufficiency.

6. Looking Forward

I think the counterexamples presented in the last two sections cause serious trouble for ARP. Perhaps some descendant of ARP can escape these counterexamples, but I have no idea what such an account would look like. What’s more, the dismal track record of counterfactual analyses in philosophy provides scant grounds for optimism on this front.¹⁵

Of course, even if we consign ARP to the graveyard of failed counterfactual analyses, that doesn’t mean that process reliabilism is dead in the water. The process reliabilist is perfectly free to jettison ARP in favour of an alternative account of defeat. But where can she find such an account?

I’ll close by briefly mentioning a potentially promising direction. In light of our counterexamples to ARP, it may be tempting to cash out defeat in terms of reasons, perhaps as follows:

Reasons Account of Defeat: S’s belief that p is defeated at t iff S has sufficiently good reason to abandon her belief that p at t .¹⁶

According to the Reasons Account of Defeat, Harry and Larry’s belief in TREE isn’t defeated since they have no good reason to abandon their belief in TREE. By contrast, Clarence’s belief that there will be a job opening is defeated, since Victor’s testimony provides him with a good reason to abandon this belief.

Given its ability to explain our counterexamples to ARP, the Reasons Account of Defeat has considerable appeal. But process reliabilists who are attracted to this account face an obvious hurdle. The Reasons Account of Defeat employs the phrase “having sufficiently good reason” in the *analysans*. Presumably, the reasons in question are *epistemic* reasons. (Or, at the very least, they include epistemic reasons.) On pain of abandoning their reductive aspirations, process reliabilists will need to cash out the notion of “having good epistemic reason” in non-epistemic terms. Perhaps this can be done, but is by no means a trivial task: process reliabilists who go down this road have their work cut out for them.¹⁷

¹⁵For a general discussion of recurrent problems with conditional analyses, see Shope (1978).

¹⁶For an influential treatment of defeat in terms of reasons, see Pollock (1987, 1994, 1995, 2001).

¹⁷Some may be tempted to analyse “good epistemic reasons” in terms of “evidence.” But this only pushes the problem back a step further: how can we cash out the notion of “having evidence” in non-epistemic terms?

It’s worth noting that Goldman’s (2011) exploration of a synthesis of evidentialism and reliabilism may be viewed as an initial step towards answering this question. There, he suggests that the concept of evidence can

7. Conclusion

In the literature on process reliabilism, the issue of defeat is often given short shrift. In this paper, I've tried to correct this oversight, arguing that the phenomenon of defeat poses a significant challenge for process reliabilists. I began by explaining why process reliabilists need a story about defeat (§2); I then argued that the standard process reliabilist story about defeat (ARP) is subject to counterexample (§4-5). It remains to be seen whether process reliabilists can give an alternative reductive account of defeat that fares better.

Of course, if it turns out that no viable reductive account of defeat is forthcoming, then this may not be a problem for process reliabilists alone. Arguably, it will spell doom for any attempt to give a reductive "Two Step" account of justification – that is, an account of justification that consists of a *prima facie* justification condition and a separate "No Defeaters" condition. If this turns out to be the case, we'll either have to give a reductive One Step theory of justification (i.e. a theory that avoids appealing to a "No Defeaters" condition), or we'll have to abandon our reductive ambitions entirely.¹⁸

References

- Kelly Becker. Epistemic luck and the generality problem. *Philosophical Studies*, 139:353–366, 2008.
- Matthew Bedke. Developmental process reliabilism: on justification, defeat, and evidence. *Erkenntnis*, 73(1):1–17, 2010.
- Laurence Bonjour. *The Structure of Empirical Knowledge*. Harvard University Press, Cambridge, MA, 1985.
- Earl Conee and Richard Feldman. The generality problem for reliabilism. *Philosophical Studies*, 89(1):1–29, 1998.
- Allan Gibbard. Rational credence and the value of truth. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, volume 2. Oxford University Press, New York, 2007.
- Alvin Goldman. What is justified belief? In Pappas, editor, *Justification and Knowledge*. Reidel, Dordrecht, 1979.

be understood in terms of reliable indication: e is evidence for p iff e is a reliable indicator of p . But even if this suggestion is right, it doesn't yet amount to an analysis of *having* evidence. After all, according to Goldman's proposal, the number of rings on a tree is evidence for a tree's age – no one needs to possess this evidence. So some further account of the "having" relation needs to be supplied, and it's not immediately obvious how to develop such an account without relying on epistemic notions.

¹⁸Thanks to David Black, Andy Egan, Alvin Goldman, Simon Goldstein, Martin J'onsson, Stephanie Leary, Lisa Miracchi, Pamela Robinson, Jonathan Schaffer, Kurt Sylvan, and two anonymous referees for helpful comments and discussion.

- Alvin Goldman. *Epistemology and Cognition*. Harvard University Press, Cambridge, MA, 1986.
- Alvin Goldman. Epistemic folkways and scientific epistemology. In *Liasons: Philosophy Meets the Cognitive and Social Sciences*. MIT Press, Cambridge, MA, 1992.
- Alvin Goldman. *Knowledge in a Social World*. Oxford University Press, New York, 1999.
- Alvin Goldman. Epistemic relativism and reasonable disagreement. In Feldman and Warfield, editors, *Disagreement*, pages 187–215. Oxford University Press, New York, 2010.
- Alvin Goldman. Towards a synthesis of reliabilism and evidentialism. In Dougherty, editor, *Evidentialism and its Discontents*, pages 254–280. Oxford University Press, New York, 2011.
- Alvin Goldman and Moshe Shaked. An economic model of scientific activity and truth acquisition. *Philosophical Studies*, 63(1):31–55, 1991.
- Thomas Grundmann. Reliabilism and the problem of defeaters. *Grazer Philosophische Studien*, 79(1):65–76, 2009.
- Michael Huemer. Compassionate phenomenal conservatism. *Philosophy and Phenomenological Research*, 74:30–55, 2007.
- Martin Jönsson. A reliabilism built on cognitive convergence: An empirically grounded solution to the generality problem. *Episteme*, 10(3):241–268, 2013.
- Martin Jönsson. Linguistic convergence in verbs for belief-forming processes. *Philosophical Psychology*, forthcoming.
- James Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4): 575–603, 1998.
- Hilary Kornblith. Justified belief and epistemically responsible action. *The Philosophical Review*, 92(1):33–48, 1983.
- Jonathan Kvanvig. Two approaches to epistemic defeat. In Baker, editor, *Alvin Plantinga*. Cambridge University Press, Cambridge, 2007.
- Keith Lehrer. *Theory of Knowledge*. Routledge, New York, 1990.
- Jack Lyons. *Perception and Basic Beliefs*. Oxford University Press, New York, 2009.
- Jack Lyons. Goldman on evidence. In Kornblith and McLaughlin, editors, *Goldman and his Critics*. Blackwell, Oxford, forthcoming.
- Sarah Moss. Scoring rules and epistemic compromise. *Mind*, 120(480):1053–1069, 2011.

Erik J. Olsson. A naturalistic approach to the generality problem. In Kornblith and McLaughlin, editors, *Goldman and his Critics*. Blackwell, Oxford, forthcoming.

John Pollock. Defeasible reasoning. *Cognitive Science*, 11:481–518, 1987.

John Pollock. Justification and defeat. *Artificial Intelligence*, 67:377–408, 1994.

John Pollock. *Cognitive Carpentry*. MIT Press, Cambridge, MA, 1995.

John Pollock. Defeasible reasoning with variable degrees of justification. *Artificial Intelligence*, 133(2):233–282, 2001.

Robert K. Shope. The conditional fallacy in contemporary philosophy. *The Journal of Philosophy*, 75(8):397–413, 1978.