

RELIABILISM AND THE SUSPENSION OF BELIEF

Weng Hong Tang

What are the conditions under which suspension of belief—or *suspension* for short—is justified? Process reliabilists hold that our beliefs are justified if and only if they are produced or sustained by reliable cognitive processes. But they have said relatively little about suspension. Perhaps they think we may easily extend an account of justified belief to deal with justified suspension. But it's not immediately clear how we may do so. As Goldman [2002: 59] notes, 'reliability is a truth-linked measure that gives positive weight to true belief and negative weight to false belief, but it ignores suspensions of judgment'. Also, Feldman and Conee [2005] claim that reliabilism, unlike evidentialism, is not well-suited to accounting for justified suspension. If they're right, evidentialism has a distinct advantage over reliabilism. Process reliabilists thus face the challenge of showing that they can account for justified suspension. In this paper, I consider some proposals on how to do so. Though several of them do not work, there are two that are promising. The first such proposal appeals to the notion of *propositional justification*; the second involves weaving *evidentialist* elements into reliabilism. I'll argue that the second proposal is better than the first.

Keywords: suspension of belief, suspension of judgement, justification, process reliabilism, indicator reliabilism, evidentialism

1. Introduction

The tallest building in the world has an odd number of windows. True or false? If you're like me, you'll be hesitant about answering either way. For lacking the relevant evidence, you'll neither believe nor disbelieve the relevant proposition. Instead, you'll suspend belief in it—your attitude towards it will be one of agnosticism.

Such suspension of belief—or suspension for short—seems justified. But suppose we're overly cautious and inclined to suspend belief in anything whose truth we can't establish with Cartesian certainty. As a result, we suspend belief about whether it'll snow despite having very good (but defeasible) evidence that it won't. Normally, such suspension is unjustified.

What are the conditions under which suspension is justified? Process reliabilists hold that our beliefs are justified just in case the cognitive processes or mechanisms that produced or sustained them are reliable. But they've said relatively little about *suspension*. Perhaps they think we may easily extend an account of justified belief to deal with justified suspension. However, Feldman and Conee [2005: 106. Italics mine] claim that though '[i]t is not possible to prove that there is nothing good for reliabilists to say about justified suspension of judgment [or belief], . . . it is clear that the theory [of reliabilism] is not well-suited to account for it'.

We'll see in a bit why one might hold such a claim. For now, note that according to Feldman and Conee, *evidentialism*—one of process reliabilism's main rivals—can handle justified suspension easily (ibid.). Evidentialists hold we may either have a justified belief that p , a justified belief that $\sim p$, or a justified attitude of suspension towards p , depending on whether the relevant doxastic attitude *fits* the evidence on which it's based [Feldman and

Conee 1985: 15, 24].¹ When we've strong enough evidence for p , such evidence will fit a belief that p . When we've strong enough evidence against p , such evidence will fit a belief that $\sim p$ (or a disbelief that p). And when our evidence neither favours p nor $\sim p$, such evidence will fit a suspended belief in p . If Feldman and Conee are right, and process reliabilism is indeed not suited to account for justified suspension, that will give evidentialism a distinct advantage over it.

Herein lies a challenge to reliabilists: show that process reliabilism—or some suitably modified version of it—has the resources to account for justified suspension. In what follows, after saying a bit more about suspension, I'll explore some attempts to meet the challenge. I'll first consider some attempts I'll argue are problematic. I'll then consider two other proposals that hold more promise. The first proposal appeals to the notion of *propositional justification*; the second weaves certain *evidentialist* elements into reliabilism. I'll argue that the second proposal is preferable to the first.

2. Reliabilism and Suspension: Some Problematic Proposals

In claiming that reliabilism isn't well-suited to accounting for justified suspension, Feldman and Conee take suspension to be an attitude. They write:

It is not just belief that can be readily evaluated by evidentialist standards. Disbelief and suspension of judgment are as readily evaluated. The justification of each *attitude* emerges in a unified and natural way from the support that the evidence provides.

[Feldman and Conee 2005: 106. Italics mine]

Similarly, Friedman [2013b: 167] holds that to suspend belief in p is to adopt an attitude of agnosticism towards p , where such an attitude 'represents (or expresses or just is) a subject's neutrality or indecision with respect to [p 's] truth'. Understood as such, suspension involves more than the mere lack of an attitude of belief or disbelief. As Wedgwood [2002: 272. Italics Wedgwood's] puts it, 'the property of neither believing nor disbelieving p is not a mental state at all—even rocks and numbers have *that* property'.

Consider again the proposition that the tallest building in the world has an odd number of windows. Someone who can't grasp the proposition or who has never considered it before may adopt no attitude towards it whatsoever. But being able to grasp the proposition and having considered it, I do have an attitude towards it, namely, one of agnosticism. Upon weighing my evidence, I find myself ruling out only some doxastic possibilities in which the proposition is false and only some doxastic possibilities in which it's true.² Since my evidence leaves it open that the proposition is false and leaves it open that it's true, my attitude towards it is one of neutrality or indecision with respect to its truth. Such neutrality or indecision does not amount to having no attitude towards the proposition—otherwise, I would not be ruling out possibilities in which it's true or false in the first place.

To be clear, I grant there may be some sense in which we're agnostic about a proposition when we adopt no attitude towards it. But I'll set such agnosticism aside. For I take it that the challenge that Feldman and Conee pose for reliabilists has bite only if suspension is an attitude. The worry for reliabilists is that, unlike evidentialists, they are unable to give a complete or unified account of the justifiedness of our doxastic attitudes. Though reliabilists may try to meet the challenge by denying that suspension is an attitude,

¹ Some philosophers (e.g. Comesaña [2010: 574–6]), have questioned whether evidentialists can provide an illuminating account of fit. For the purposes of this paper, I'll grant that they can do so.

² For example, I can at least rule out the possibility in which the tallest building in the world has an even number of windows in virtue of having exactly two windows, as well as the possibility that it has an odd number of windows in virtue of having exactly three windows.

such a denial would at best be controversial. It's better for reliabilists to show that suspension poses no problems even if it's an attitude.

At this point, one may think that if suspension is an attitude, it amounts to having a credence or degree of confidence that is neither high enough to qualify as a belief nor low enough to qualify as a disbelief. In such a case, a reliabilist account of justified credence will yield a reliabilist account of justified suspension. One may thus wonder if we should focus in the first instance on the former.³

The project of giving a reliabilist account of justified credence is a worthy one. Later, I'll say more about reliabilism and justified credences. But the issue of whether suspensions are reducible to credences is contentious. For instance, Friedman [2013b: 180] suggests that 'the agnostic attitude . . . is sui generis'. She also thinks we may sometimes adopt an attitude of suspension towards propositions to which we assign no credence whatsoever.⁴ Taking suspension to be an attitude, she writes:

A subject might be so utterly in the dark about . . . whether the Hill 50 Gold Mine was Australia's most profitable mine between 1955 and 1961 that he ought to simply refuse to have any degrees of belief [in the relevant proposition]. But it is epistemically permissible that he suspend judgment about [it] even if we think he can't or shouldn't assign credences to [it].

[Friedman 2013a: 66]

Of course, Friedman may prove to be wrong. Perhaps, if we can't assign any credence to a proposition, then we simply have no attitude—not even one of agnosticism—towards it. Perhaps, insofar as suspension is an attitude, it involves nothing more than the assignment of middling credences. Such issues are controversial, and I can't settle them here. But ideally, a reliabilist account of doxastic attitudes should not be held hostage to which view of suspension eventually proves to be correct. Precisely because the issues are controversial, it's worth coming up with a reliabilist account of justified suspension that is neutral with respect to them. In the rest of this section, I'll consider some proposals.

According to process reliabilism, a belief is justified if and only if it's produced (or sustained) by a *reliable* cognitive process or system of cognitive processes, that is, one that tends to yield a high proportion of true to false beliefs [Goldman: 1979]. For example, whereas careful observation tends to yield a high proportion of true beliefs, wishful thinking tends not to do so. So, according to reliabilists, if I believe there's a table in front of me as a result of observing my environment carefully, my belief is caused by a reliable process and therefore justified. If I believe I'll live forever as a result of wishful thinking, my belief is caused by an unreliable process and therefore unjustified.

But how may reliabilists account for whether a suspended belief is justified? One may suggest that if suspensions are produced by cognitive processes that also produce beliefs, a suspended belief is justified if and only if the process that produced it is reliable. However, the suggestion doesn't work. Let's grant that suspensions are produced by processes that also produce beliefs.⁵ The reliability of such processes is independent of whether the suspensions

³ Thanks to an anonymous referee for pressing me on this point and for pressing me to clarify whether I take suspension to be an attitude.

⁴ See Decker [2012] too (though Decker considers a case in which we have no attitude whatsoever towards *p* as one in which we suspend belief in *p*).

⁵ As Feldman and Conee [2005] observe, a process reliabilist such as Goldman will be happy to grant this. In the latter's own words,

If *S*'s belief in *p* at *t* results from a reliable cognitive process, and there is no reliable or conditionally reliable process available to *S* which, had it 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 in *p* at *t* is justified.

[Goldman 1979: 20]

they produce are justified. For *suspensions are not taken into account when computing the ratio of true to false beliefs*. Or as Goldman [2002: 59] puts it, ‘reliability is a truth-linked measure that gives positive weight to true belief and negative weight to false belief, but it ignores suspensions of judgment’.

Two examples, the first of which shows that reliability is not sufficient for justified suspension, and the second of which shows that it's not necessary, will reinforce the point that the suggestion above does not work. First, consider Sceptical Sue, whose overly conservative cognitive processes result in suspension unless her evidence *entails* that the relevant propositions are true (or false).⁶ With such exacting standards, the (relatively few) beliefs the processes produce may mostly be true. In that case, the processes are reliable and, by the lights of the process reliabilist, the beliefs they produce justified. But this shows that reliability isn't sufficient for justified suspension. For there'll be many cases in which Sceptical Sue has very strong inductive or abductive evidence for believing propositions not entailed by such evidence. In such cases, suspending belief in the relevant propositions seems unjustified. But that is exactly what Sceptical Sue's cognitive processes will lead her to do.

Next, consider Credulous Callum, whose cognitive processes lead to suspension in p whenever the evidence for p is weighted equally against the evidence for $\sim p$. Such suspension may well be justified. But in such a case, must those cognitive processes also produce beliefs the majority of which are true? No, for Credulous Callum may live up to his name too well—the same cognitive processes may produce a belief that p whenever the evidence for p just slightly outweighs that for $\sim p$. Since such processes are likely to produce a significant proportion of false beliefs, reliability isn't necessary for justified suspension.

It doesn't help to suggest that a process is reliable just in case the proportion of true beliefs among all the doxastic states it produces, be they states of belief or suspension, is high. For this has the implausible consequence that there would be *no* epistemic difference between suspending belief and having false beliefs: either would lead to an *equal* decrease in the proportion of true beliefs. For example, suppose that T is a set of ten true propositions and F a set of ten false propositions. Consider a process that leads one to believe all the propositions in T and in F , and another process that leads one to believe all the propositions in T but suspend belief in all those in F . On the current suggestion, the doxastic states produced by the first process are as justified as those produced by the second. But all things being equal, believing the false propositions in F seems less justified than suspending belief in them.

It also doesn't help to hold that a suspended belief is justified if and only if the process that produced it produces neither a high nor a low ratio of true beliefs. For a reliable process may sometimes lead to justified suspension. Consider the process of forming beliefs based on careful observation. Such a process will presumably produce a high ratio of true to false beliefs. But it may sometimes lead to suspension instead of belief. When carefully observing a flying object in the distance, we may have a visual experience as of a bird, though not a visual experience as of any particular species of bird. This may allow us to form a justified belief that we are looking at a bird, but not a justified belief about whether it is, say, a magpie. In such a case, the relevant process may lead to justified suspension about the

As Feldman and Conee [2005: 105] note, ‘Goldman proposes that the justification of a belief is defeated when one has available a reliable belief-forming process that would lead to not having that belief’. Of course, one way a process may lead to not having a belief that p is by leading to a disbelief that p . But it's reasonable to think that another way is by leading to suspension in p .

⁶ For example, suppose Sceptical Sue has a veridical perception of an apple on the table. Her cognitive processes may take her perceptual experience as input and produce as output a true belief in the proposition ‘There *seems* to be an apple on the table’, but at the same time, lead to suspension in the proposition ‘There *is* an apple on the table’. For her having the perceptual experience she has may entail the first proposition but not the second.

species of the bird observed.⁷

Further, it won't do to propose simply that a suspended belief produced by a process is justified if and only if neither a low nor a high proportion of all the cases in which the process leads to suspension are ones in which the relevant propositions are true. The following example from Feldman and Conee [2005: 106] shows why. A coin we know to be symmetrical is tossed repeatedly. After every toss, we get to see which side faces up. Suppose there's a process that, time and again, leads us to suspend belief in the coin landing heads—after seeing it land. Since it's likely that half of the cases in which we suspend belief in this manner are ones in which the coin lands heads, such suspension is justified given the current proposal. But as Feldman and Conee [ibid.: 16] point out, such suspension is *not* justified.⁸

I'll consider one more proposal that seems more promising than those considered so far but that I'll argue is still problematic. Recall that, according to Goldman [2002: 59], 'reliability is a truth-linked measure that gives positive weight to true belief and negative weight to false belief, but it ignores suspensions of judgment'. Perhaps, to account for justified suspension, we should look for a measure that gives some weight to suspensions. And one way to do so is to appeal to some sort of scoring rule.

Joyce [1998] has proposed that we appeal to Brier scoring to measure the accuracy of *credences* (of which more later). The Brier score of a credence of x in p is calculated with the formula $(x - T(p))^2$, where $T(p)$ equals 1 if p is true and 0 if p is false.⁹ The lower the Brier score of one's credence, the more accurate it is. For example, suppose p is true. Then a credence of 0.7 in p is more accurate than a credence of 0.4 in p since the Brier score of the former is 0.09 whereas that of the latter is 0.36. In the extreme, a credence of 1 in p is perfectly accurate whereas a credence of 0 in p is perfectly inaccurate.

Though Brier scoring gives us a fine-grained measure of the accuracy of credences, one may adapt it for a more coarse-grained measure of the accuracy of beliefs, disbeliefs, and suspensions. Suppose we treat a belief as we would a credence of 1, a disbelief as we would a credence of 0, and a suspension as we would a credence of 0.5. Then a true belief will yield a Brier score of 0, a false belief a score of 1, and a suspended belief (in either a truth or a falsehood) a score of 0.25. Appealing to Brier scoring, we may formulate the following theory:

(Brier) A subject's doxastic attitude towards p (be it an attitude of belief, disbelief, or suspension) is justified if and only if it's caused by a process that tends to result in doxastic attitudes (beliefs, disbeliefs, or suspensions) that have a low average Brier score (i.e. that are on average highly accurate).

By cashing out reliability in terms of the ratio of true to false beliefs, process reliabilism comports with the view that there is something epistemically good about believing truths and

⁷ Also, as seen in footnote 5, Goldman should have no objection to holding that a reliable cognitive process may lead to suspension. Presumably, he'd also hold that it may lead to justified suspension.

⁸ The proposal faces another problem. Given that a proposition is either true or false, and that any process leading to suspension in a proposition will also lead to suspension in its negation, exactly half of the propositions in which a process leads to suspension will be true. Hence, the proposal seems to make it *too easy* for a process to lead to justified suspension.

I'll not dwell too much on whether the problem can be solved, since the objection raised by Feldman and Conee above pose a big enough worry for the proposal. Perhaps one may maintain that a state of suspension in p is really the same as a state of suspension in $\sim p$. One may then suggest that when we consider cases in which a process leads to suspension, we should exclude cases in which we suspend belief in propositions of the form $\sim p$ in order to avoid double-counting. Whether this solves the problem depends on whether there's a principled, non-arbitrary way to determine whether a proposition has the form p or $\sim p$; otherwise, there'll be no non-arbitrary answer to the question 'What is the proportion of true propositions in which one suspends belief?'

⁹ Brier scoring was first put forward by Brier [1950] to gauge the accuracy of weather forecasts.

not believing falsehoods. But focusing only on this ratio leaves us unable to account for suspension. (Brier) aims to remedy this problem—by appealing to accuracy, it aims to capture the idea that, epistemically speaking, believing a truth is better than suspending belief, which is in turn better than believing a falsehood.¹⁰

(Brier) makes room for suspension without straying from the spirit of process reliabilism. It says that justified doxastic attitudes are produced by a certain kind of *process*. It also maintains that this kind of process is in some sense *truth-conducive*, the main difference between traditional process reliabilism and (Brier) being that whereas the former takes a truth-conducive process to be one that produces a high ratio of *true* beliefs, the latter takes it to be one that produces doxastic attitudes that are on average highly *accurate*.

But (Brier) faces the following worry. It says that whether a suspended belief is justified depends on whether the process responsible for it is associated with a low average Brier score. This in turn depends on whether the process produces, on average, more true beliefs than false beliefs and suspensions. But though the measure employed by (Brier), unlike the measure of reliability, does give some weight to suspension, it still ties the forming of justified suspensions *too closely* to the forming of justified beliefs. Our cognitive processes may produce justified beliefs while producing many unjustified suspensions. And they may produce justified suspensions while producing many unjustified beliefs. But (Brier) rules out such possibilities *in principle*.

Recall Sceptical Sue, whose overly conservative cognitive processes result in suspension unless her evidence entails that the relevant propositions are true (or false). As we've seen, many of her suspended beliefs may be unjustified—she may suspend belief in many true propositions we're ordinarily justified in believing on the basis of induction or abduction. But the few beliefs her cognitive processes produce may mostly be true and justified. Such a case is possible, but (Brier) says otherwise. For Sceptical Sue's suspending belief in many true propositions will lead to her doxastic states having a (relatively) high average Brier score, which means, by the lights of (Brier), that her beliefs are unjustified. So much the worse for (Brier).

Also, recall Credulous Callum, whose cognitive processes tend to result in suspension in *p* whenever the evidence for *p* is weighted equally against the evidence for $\sim p$, but too readily produce a belief that *p* whenever the evidence for *p* just slightly outweighs that for $\sim p$. In such a case, Credulous Callum's suspensions may well be justified even if he's also likely to have many false and unjustified beliefs. Such a case is possible, but (Brier) says otherwise. For given Credulous Callum's many false beliefs, the average Brier score of his doxastic states will tend to be high, which means, by the lights of (Brier), that his suspensions are unjustified. Again, so much the worse for (Brier).

While my focus is on suspension and I do not assume that suspensions are reducible to credences, one may wonder if the problems above arise only because of the ham-fistedness of a proposal couched in coarse-grained terms (as an anonymous referee puts the point). To answer the worry, it's worth noting that the problems that Sceptical Sue and Credulous Callum pose for (Brier) remain even if to suspend belief in *p* is to have an intermediate

¹⁰ One may get theories similar to (Brier) by appealing to other scoring rules. Following Goldman [2002: 58. *Italics mine*], we may hold that 'believing a truth carries more *veritistic value* than suspension of judgment; and suspension of judgment carries more veritistic value than disbelief'. Now, suppose the veritistic value of a true belief is represented by a score of 1, that of a false belief by a score of 0, and that of a suspended belief by a score of 0.5 [Goldman 1999: 89]. One may then put forward the following theory:

(Verity) A subject's doxastic attitude towards *p* (be it an attitude of belief, disbelief, or suspension) is justified if and only if it's caused by a process that tends to result in doxastic attitudes (beliefs, disbeliefs, or suspensions) that have a high average veritistic value.

For brevity's sake, I'll not discuss such a theory further. But what I say about (Brier) in what follows should apply, *mutatis mutandis*, to (Verity).

credence of, say, 0.5 in p .

Someone who wishes to appeal to Brier scoring to account for justified credences may suggest the following theory:

(Brier-C) A subject's credence in p is justified if and only if it's caused by a process that tends to result in credences that have a low average Brier score. (Cf. Lam [2011: 215–16].)

But consider Sceptical Sue, whose overly conservative cognitive processes result in middling credences unless her evidence entails that the relevant propositions are true (or false). Many of such credences may be unjustified. Sceptical Sue may assign middling credences to many true propositions to which we're ordinarily justified in assigning high credences on the basis of induction or abduction. But the few high credences her cognitive processes produce may mostly have true contents and be justified. Such a case is possible, but (Brier-C) says otherwise. For Sceptical Sue's assignment of middling credences to many true propositions will lead to her doxastic states having a (relatively) high average Brier score. By the lights of (Brier-C), this means that even the high credences produced by the relevant cognitive process are unjustified.

Also, consider Credulous Callum, whose cognitive processes tend to result in a credence of 0.5 in p whenever the evidence for p is weighted equally against the evidence for $\sim p$, but too readily produce a very high credence in p whenever the evidence for p just slightly outweighs that for $\sim p$. In such a case, Credulous Callum's credences of 0.5 may well be justified even if he's also likely to assign unjustified high credences to several false propositions. Such a case is possible, but (Brier-C) says otherwise. For given that Credulous Callum assigns very high credences to a substantial number of false propositions, the average Brier score of his credences will tend to be high. By the lights of (Brier-C), this means that even his credences of 0.5 are unjustified.

3. Two More Proposals: (Propositional) and (Evidence)

We've seen various failed attempts to account for justified suspension on behalf of reliabilists. This might make us think that reliabilism is just 'not well-suited' to the task [Feldman and Conee 2005: 106]. But reliabilists shouldn't capitulate: their pool of resources is deeper than might first appear.

As we've seen, evidentialists hold that suspension in p is justified only if it fits the evidence on which it's based. And presumably, evidence for p fits an attitude of suspension towards p just in case such evidence neither fits a belief that p nor a belief that $\sim p$. Question: is there a reliabilist analogue of fit? If so, perhaps reliabilists may exploit it to account for justified suspension. In what follows, I'll consider two different ways to spell out such an analogue.

3.1 *An Appeal to Propositional Justification*

Reliabilists are typically concerned with *doxastic* justification, which has to do with the conditions under which a belief is justified. But they have the means to account for *propositional justification*, which has to do with the conditions under which there is justification for a belief, whether or not we have the belief.¹¹ Goldman [1979: 21], for instance, writes:

Person S is [propositionally] justified in believing p at [time] t if and only if there is a reliable belief-forming operation available to S which is such that if S had applied that operation to his total cognitive state at t , S would believe p at t -plus-

¹¹ For more on the distinction between doxastic and propositional justification, see Firth [1978].

delta (for a suitably small delta) and that belief would be [doxastically] justified.

By appealing to propositional justification, understood in reliabilist terms, we can cash out a reliabilist analogue of fit and exploit it to account for justified suspension. Whereas evidentialists ask whether one's belief fits the relevant evidence, reliabilists may ask whether there's propositional justification for one's belief. And whereas evidentialists hold that suspension in p is justified only if one's evidence fits neither a belief nor a disbelief that p , reliabilists may hold that suspension in p is justified only if one lacks propositional justification for believing p and for believing $\sim p$. As a first pass, consider:

(Propositional- α) S justifiably suspends belief in p just in case

1. S suspends belief in p ;
2. S lacks propositional justification for believing p ; and
3. S lacks propositional justification for believing $\sim p$.¹²

Suppose propositional justification can be cashed out in terms of available and reliable cognitive processes in the way Goldman [1979] suggests. If, given (Propositional- α), we can cash out justified suspension in terms of the lack of propositional justification, we'll also be able to cash it out in terms of reliable cognitive processes—or the lack thereof.

It's easy to verify that (Propositional- α) avoids the problems faced by the proposals discussed earlier. For instance, (Brier) maintains that justified suspensions can only be produced by a process that tends to produce a high proportion of true beliefs. (Propositional- α), to its credit, imposes no such implausible condition. Also, according to an earlier proposal, a suspended belief produced by a process is justified if and only if neither a low nor a high proportion of all the cases in which the process leads to suspension are ones in which the relevant propositions are true. As we've seen, this proposal yields the wrong result in a case in which we suspend belief about whether a symmetrical coin has landed heads after seeing it land. But (Propositional- α) faces no such problem. After each toss of the coin, a reliable cognitive process is made available to us—upon having a visual experience of the coin's landing, we gain propositional justification either for believing that it has landed heads or for believing that it hasn't. In each case, by the lights of (Propositional- α), suspending belief about whether the coin has landed heads is unjustified.

But (Propositional- α) faces a problem of its own. It's well known that we may have propositional justification for believing p and yet have an unjustified belief that p (e.g. see Pollock and Cruz [1999: 35–6]). This happens when our belief that p is formed for the wrong reasons—or, to speak in reliabilist terms, formed via an unreliable process even though a reliable one is available. So an account according to which we justifiably believe p just in case we believe p and have propositional justification for doing so faces a serious problem. (Propositional- α) succumbs to a similar problem. Suppose we lack propositional justification for believing p and for believing $\sim p$, but we suspend belief in p only because we are too lazy to examine our evidence to figure out what to believe (for instance). Our suspension is unjustified. But (Propositional- α) says otherwise.

To solve the problem, the process reliabilist may refine (Propositional- α) to get:

(Propositional) S justifiably suspends belief in p just in case

1. S suspends belief in p ;
2. S lacks propositional justification for believing p ;

¹² I thank Ben Blumson, Mark D'Cruz, and Michael Pelczar for suggesting versions of (Propositional- α).

3. S lacks propositional justification for believing $\sim p$; and
4. S 's suspension in p is formed by a process which tends to be such that most of the propositions in which it leads to suspensions are ones that S lacks propositional justification for believing and for disbelieving.

(Propositional) avoids the problem that befalls (Propositional- α). Suppose we lack propositional justification for believing p and for believing $\sim p$, but we suspend belief in p only because we're too lazy to examine our evidence to figure out what to believe. This way of suspending belief will tend to lead to suspension in a substantial number of cases in which we've propositional justification for either believing or disbelieving the propositions in question. So clause 4 of (Propositional) isn't satisfied; accordingly, our suspension in p is deemed unjustified.

3.2 *An Appeal to Evidence*

(Propositional) offers us a reliabilist account of justified suspension via a reliabilist analogue of fit. But let's now turn to a different proposal inspired by Alston. According to Alston [2005: 99], to have a justified belief that p is to have a belief that p based on an *adequate* ground, where such a ground is adequate just in case the objective probability of p being true given that the belief is based on that ground is very high (and where *objective probability* is cashed out in terms of hypothetical relative frequencies).¹³ Correspondingly, we may hold that to have a justified disbelief that p is to have a disbelief based on a certain ground, where the objective probability of p being true given that the disbelief is based on that ground is very low. Speaking in terms of fit, our belief (or disbelief) that p fits the ground on which it's based just in case the relevant objective probability is very high (or low).

But what is it for a doxastic attitude to be *based* on a certain *ground*? Suppose we've a visual experience as of a tree in the yard, and this experience serves as the input to a cognitive process that produces the belief that there's a tree in the yard. Then our experience is the ground on which the belief in question is based. More generally, a doxastic attitude is based on a certain ground just in case that ground serves as the input to the cognitive process responsible for producing the attitude in question. Now, the ground of a doxastic attitude may also be thought of as one's evidence for that attitude. But to avoid circularity, we shouldn't cash out *ground* or *evidence* as that which justifies one's doxastic attitudes.¹⁴ And Alston [ibid.: 83] doesn't; he holds that a ground is 'something psychological—some psychological state or process' such as a belief, a memory, or an experience. For example, our belief that there's a tree in the yard is based on something psychological, namely, a visual experience.

Though Alston [2005] focuses on belief, it's natural, given the preceding, to hold that our suspension in p fits the ground on which it's based just in case the relevant objective probability is neither very high nor very low. Correspondingly, it's natural to suggest, at least as a first pass, the following account of justified suspension:

(Evidence- α) S justifiably suspends belief in p just in case

1. S suspends belief in p based on some ground g , and
2. the objective probability of p being true given that S 's suspension is based on g

¹³ Strictly speaking, Alston [2005: ch. 1] (unlike Alston [1988]) isn't concerned with defending a theory of doxastic justification. But what he says about reliability will still help us formulate a reliabilist theory of suspension.

¹⁴ Alston [2005: 82–3] prefers the term 'ground' to the term 'evidence'. For my purposes, it's fine to take them to be synonymous.

equals some value x that is neither very high nor very low.¹⁵

At this point, one may wonder if (Evidence- α) is a kind of process reliabilism. After all, Alston's theory of justified belief is usually thought to be a kind of *indicator* reliabilism, according to which the justifiedness of a belief depends on whether the belief's being based on a certain ground is 'sufficiently indicative of [its] truth' [Alston 1988: 269]. But Alston [2005: 136–7] has recently argued that his theory *is* a kind of process reliabilism. As mentioned, the ground on which a belief is based is the input to the cognitive process that has the belief as an output. And according to Alston, a belief b 's being based on ground g is sufficiently indicative of its truth just in case the objective probability of b being true given that it is based on g is high, i.e., just in case the objective probability of b being true given that b is the output of a cognitive process with g as the input is high. Since Alston takes the relevant kind of objective probability to be relative frequency, this means that b 's being based on g is sufficiently indicative of its truth just in case the frequency of true belief outputs relative to the relevant cases—cases similar to that in which the process takes g as its input and produces b as its output—is high. But this means the cognitive process is reliable (with respect to the relevant cases). As Alston writes, 'reliability of process and reliability of indicator turn out to coincide' [ibid.: 137].

We may interpret (Evidence- α) along similar lines. Suppose that s , a suspended belief in p that is based on g , is justified. Then, according to (Evidence- α), the objective probability of p being true given that s is based on g is neither low nor high. Taking 'objective probability' to refer to relative frequency, we may read the preceding as saying that the process that led to s produces a middling frequency of suspensions in true propositions relative to the relevant cases—cases similar to that in which the process takes g as its input and produces s as its output. Interpreted as such, (Evidence- α) keeps to the spirit of process reliabilism. First, it says that justified suspensions are produced by a certain kind of *process*. Second, like traditional process reliabilism, (Evidence- α) cashes out justified doxastic states in terms of *frequency of truth*: while the former ties justified belief to a high frequency of true beliefs, the latter ties justified suspension to a middling frequency of suspensions in true propositions.

But (Evidence- α) faces a problem similar to that faced by an earlier proposal. Suppose that, based only on our visual experience of a symmetrical coin being tossed, we suspend belief about its landing heads. If the coin is fair, it's likely that its frequency of landing heads—relative to cases similar to that in which we suspend belief about the coin landing heads based only on such a kind of experience—is neither high nor low. Then, by the lights of (Evidence- α), our suspension is justified. But it *isn't* justified if we've neglected certain evidence, say, if we've also had a visual experience of the coin landing heads but ignored it.

To deal with the problem, let's modify (Evidence- α) to get:

(Evidence) S justifiably suspends belief in p just in case

1. S suspends belief in p based on some ground g ;
2. the objective probability of p being true given that S 's suspension is based on g equals some value x that is neither very high nor very low; and
3. there is no more inclusive ground g' had by S such that the objective probability of p being true given that S 's suspension is based on g' does not equal x .¹⁶

¹⁵ For simplicity, I'll focus on cases involving non-doxastic grounds. To deal with cases involving doxastic grounds, we may add a clause to (Evidence- α) to the effect that if g consists of doxastic states, then those states are themselves justified.

¹⁶ Condition 3 of (Evidence) is inspired by Feldman and Conee's [1985: 24] account of doxastic justification

(Evidence) solves the problem above. In the coin example, our suspension of belief based only on our experience of a symmetrical coin being tossed is unjustified. For there's a more inclusive ground—namely, one involving both this experience and an experience of the coin landing heads—such that the objective probability of the coin landing heads, given that our suspension is based on that more inclusive ground, is very high.

It's also easy to verify that (Evidence) avoids the problems faced by the proposals considered in section 2. For example, (Brier) is committed to there being a very strong link between justified suspension and reliability. But according to (Evidence), for the purposes of evaluating whether a suspended belief produced by a cognitive process is justified, we should focus only on cases in which the process produces suspensions. And for the purposes of evaluating whether a belief produced by a cognitive process is justified, we should focus only on cases in which the process produces beliefs. Hence, (Evidence), unlike (Brier), avoids positing too strong a link between justified suspension and reliability.

4. (Evidence) vs. (Propositional)

We've seen that process reliabilists have the resources to account for justified suspension. In fact, there are at least two natural ways of doing so. I'll argue, however, that (Evidence) is preferable to (Propositional).

Here's why. Though this paper focuses on justified suspension, one may wonder if reliabilists can account for justified *credence*. After all, process reliabilism cashes out the justifiedness of beliefs in terms of the truth-conduciveness of cognitive processes. But a credence of 0.5 (say), like a suspended belief, is not the kind of thing that admits of truth or falsity. Given this, one may wonder if we can extend (Propositional) and (Evidence) to account for justified credence.

(Evidence) can be extended easily to account for justified credence. Recall that it attempts to account for justified suspension via a reliabilist notion of fit understood in terms of objective probability. According to it, our suspension in p fits the ground on which it's based just in case the objective probability of p being true, given that our suspension is based on the ground in question, is neither high nor low. Correspondingly, we may hold that S 's credence of x in p based on g is fitting just in case the objective probability of p being true given that S 's credence is based on g equals (or approximates) x . With this notion of fit in mind, it's easy to come up with an accompanying account of justified credence. Consider:

(Evidence-C) S 's credence of x in p is justified just in case

1. S assigns a credence of x to p based on some ground g ;
2. the objective probability of p being true given that S 's credence is based on g equals (or approximates) x ; and
3. there is no more inclusive ground g' had by S such that the objective probability of p being true given that S 's credence is based on g' does not equal (or

or *well-foundedness* (in particular, clause (c) of condition (ii) below). According to them: S 's doxastic attitude D at t toward proposition p is well-founded if and only if

- (i) having D toward p is justified for S at t ; and
- (ii) S has D toward p on the basis of some body of evidence e , such that
 - (a) S has e as evidence at t ;
 - (b) having D towards p fits e ; and
 - (c) there is no more inclusive body of evidence e' had by S at t such that having D towards p does not fit e' .

[ibid.]

approximate) x .¹⁷

If suspensions are reducible to credences, (Evidence) may be seen as a coarse-grained version of (Evidence-C). But if not, we may take (Evidence) and (Evidence-C) to give us a unified treatment of justified suspension and credence.¹⁸

It's hard, however, to see how we may extend (Propositional) to account for justified credence. (Propositional) understands the reliabilist analogue of fit in terms of propositional justification: our suspension in p is fitting just in case we lack propositional justification for believing p and for believing $\sim p$. But it's hard to see how to provide a similar account of fit for credences. For example, a lack of propositional justification for believing that it'll rain seems necessary for a credence of 0.4 in rain to be fitting. But it's not sufficient. Relatedly, what would explain the difference between a case in which only a credence of 0.5 in rain is justified and one in which only a credence of 0.7 in rain is justified? The notion of propositional justification, cashed out in terms of reliability, is too coarse grained to provide the requisite explanation.

All other things being equal, having a unified treatment of both justified suspension and credence seems better than having to treat them disparately. For this reason, I prefer (Evidence) to (Propositional). But one might worry that not all other things are equal—that (Evidence), by invoking the notion of basing one's doxastic attitudes on one's grounds, departs too radically from traditional process reliabilism and concedes too much to the evidentialist. For this reason, one might prefer (Propositional) to (Evidence). After all, when it comes to accounting for justified suspension, (Propositional) does not appeal to anything beyond the reliabilist's usual arsenal.

Note, however, that even process reliabilists have of late been happy to embrace a hybrid version of process reliabilism that incorporates evidentialist elements. Comesaña [2010: 385], inspired by Alston [1988], defends a version of process reliabilism that incorporates the notion of grounds or evidence; he thinks that such a hybrid version of reliabilism solves or at least alleviates various problems that plague traditional process reliabilism.¹⁹ Though Comesaña [ibid.: 582, 597] focuses on justified belief, he notes in passing that his account can be extended to deal with justified suspension. If what I say in this paper is correct, this lends even more credence to his account (though this is not a point that Comesaña argues for explicitly).

Notably, Goldman [2011: 263]—the foremost proponent of process reliabilism—has also floated the view that the incorporation of evidentialist elements into the theory is a 'salutary addition'; he notes that its 'attraction . . . seems especially obvious in inferential justification', where one doxastic state is justified on the basis of another. Focusing on credences, he suggests that if p is one's 'total doxastic evidence' for q , and one has no other non-doxastic evidence for q , then one's credence of x in q fits one's evidence for q if and only if the degree of confirmation p confers upon q is x [ibid.]. Though Goldman [ibid.] does not discuss suspension, we may illustrate his point in terms of beliefs and suspensions. Suppose we're wondering whether p and our only relevant evidence is a stock of justified

¹⁷ I discuss this account of justified credence in greater detail elsewhere [Tang forthcoming].

¹⁸ Some philosophers hold that a credence of x in p is really a binary belief that the objective probability of p being true is x (e.g. Harman [1986: 24]; Pollock [2006: 94]). On this view, there's no need for a reliabilist account of credence over and above a reliabilist account of belief. Still, a fair number of philosophers deny that credences can be reduced to binary beliefs about objective probabilities (e.g. Christensen [2004: 18–20]; Frankish [2009: 77–8]). So it's worth looking for an account of justified credence compatible with such a denial.

¹⁹ Some such problems include the generality problem (see Conee and Feldman [1998]) and the problem posed by subjects who form true beliefs reliably due to their powers of clairvoyance but whose beliefs formed in such a manner are intuitively unjustified (see BonJour [1980: 59–61]).

beliefs B . Then if the probability of p being true given that our beliefs in B are all true is neither high nor low, suspension in p seems more fitting than either belief or disbelief in p . This helps explain why, intuitively, our beliefs in B justify suspension, but neither belief nor disbelief, in p .

The appeal to Comesaña [2010] and Goldman [2011] is not a mere appeal to authority; rather, it's meant to suggest that the price of any slight departure from traditional process reliabilism is worth paying—the incorporation of evidentialist elements into process reliabilism yields various benefits, not least the provision of a unified treatment of both justified suspension and credence.

And (Evidence) involves but a *slight* departure from traditional process reliabilism. With help from Alston, we've seen that (Evidence) is a kind of process reliabilism that relates justification to frequency of truth. Admittedly, (Evidence), as well as the theories above, departs slightly from tradition by positing relations of fit between our grounds and the world. But such grounds or evidence are to be understood in terms of mental states, and as Goldman [ibid.: 263. Italics Goldman's] observes, while traditional process reliabilism does not speak of evidential fit, it 'has never hesitated to invoke mental *states* . . . in its set of resources'.²⁰ Further, in Goldman's words, the mental states that are our grounds or evidence 'qualify as items of evidence (ultimately) because they—or their ilk—stand in reliable-indicator relationships to *facts in the world*' [ibid.: 257. Italics mine]. Feldman and Conee, as Goldman notes, are 'unlikely to applaud this maneuver' [ibid.]. For it amounts to disavowing the evidentialist view that justification consists *entirely* in having certain grounds or evidence. In relating justification to frequency of truth—an externalist condition through and through—(Evidence) remains faithful to the spirit of traditional process reliabilism.

5. Conclusion

The main aim of this paper has been to show that, *pace* Feldman and Conee [2005: 106], process reliabilism is 'well-suited' to account for justified suspension. Admittedly, certain proposals on how to provide such an account are untenable. But both (Propositional) and (Evidence) fare better. Whichever of the two you prefer, justified suspension poses no problem for process reliabilists. However, (Evidence) is preferable to (Propositional). The former, unlike the latter, affords us a unified treatment of both justified suspension and credence.²¹

National University of Singapore

²⁰ For instance, when giving examples of belief-forming processes, Goldman [1979: 11–12] talks about 'a memory process, which takes as input beliefs or experiences at an earlier time and generates as output beliefs at a later time', and 'reasoning processes, where the inputs include antecedent beliefs and entertained hypotheses'.

²¹ Many thanks to Jens Christian Bjerring, Ben Blumson, Mark D'Cruz, Yongming Han, Michael Pelczar, and two anonymous referees for their invaluable comments.

References

- Alston, W.P. 1988. An Internalist Externalism, *Synthese* 74/3: 265–83.
- Alston, W.P. 2005. *Beyond “Justification”*: *Dimensions of Epistemic Evaluation*, Ithaca: Cornell University Press.
- BonJour, L. 1980. Externalist Theories of Empirical Knowledge, *Midwest Studies in Philosophy* 5/1: 53–73.
- Brier, G.W. 1950. Verification of Forecasts Expressed in Terms of Probability, *Monthly Weather Review* 78/1: 1–3.
- Christensen, D. 2004. *Putting Logic in Its Place: Formal Constraints on Rational Belief*, New York: Oxford University Press.
- Comesaña, J. 2010. Evidentialist Reliabilism, *Noûs* 44/4: 571–600.
- Conee, E. and R. Feldman 1998. The Generality Problem for Reliabilism, *Philosophical Studies* 89/1: 1–29.
- Decker, J. 2012. Disagreement, Evidence, and Agnosticism, *Synthese* 187/2: 753–83.
- Feldman, R. and E. Conee 1985. Evidentialism, *Philosophical Studies* 48/1: 15–34.
- Feldman, R. and E. Conee 2005. Some Virtues of Evidentialism, *Veritas* 50/4: 95–108.
- Firth, R. 1978. Are Epistemic Concepts Reducible to Ethical Concepts?, in *Values and Morals: Essays in Honor of William Frankena, Charles Stevenson, and Richard Brandt*, ed. A.I. Goldman and J. Kim, Dordrecht: D. Reidel Publishing Co.: 215–29.
- Frankish, K. 2009. Partial Belief and Flat-out Belief, in *Degrees of Belief*, ed. F. Huber and C. Schmidt-Petri, Dordrecht: Springer: 75–93.
- Friedman, J. 2013a. Rational Agnosticism and Degrees of Belief, in *Oxford Studies in Epistemology* Vol. 4, ed. T.S. Gendler and J. Hawthorne, Oxford: Oxford University Press.
- Friedman, J. 2013b. Suspended Judgment, *Philosophical Studies* 162/2: 165–81.
- Goldman, A.I. 1979. What is Justified Belief?, in *Justification and Knowledge: New Studies in Epistemology*, ed. G.S. Pappas, Dordrecht: D. Reidel Publishing Co.: 1–23.
- Goldman, A.I. 1999. *Knowledge in a Social World*, New York: Oxford University Press.
- Goldman, A.I. 2002. *Pathways to Knowledge: Private and Public*, New York: Oxford University Press.
- Goldman, A.I. 2011. Toward a Synthesis of Reliabilism and Evidentialism? or: Evidentialism’s Problems, Reliabilism’s Rescue Package, in *Evidentialism and its Discontents*, ed. T. Dougherty, New York: Oxford University Press: 254–80.
- Harman, G. 1986. *Change in View: Principles of Reasoning*, Cambridge, MA: The MIT Press.
- Joyce, J.M. 1998. A Nonpragmatic Vindication of Probabilism, *Philosophy of Science* 65/4: 575–603.
- Lam, B. 2011. On the Rationality of Belief-Invariance in Light of Peer Disagreement, *Philosophical Review* 120/2: 207–245.
- Pollock, J.L. 2006. *Thinking about Acting: Logical Foundations for Rational Decision Making*, New York: Oxford University Press.
- Pollock, J.L. and J. Cruz 1999. *Contemporary Theories of Knowledge*, 2nd edition, Maryland: Rowman & Littlefield Publishers, Inc.
- Tang, W.H. forthcoming. Reliability Theories of Justified Credence, *Mind*.
- Wedgwood, R. 2002. The Aim of Belief, *Philosophical Perspectives* 16: 267–97.