Pettigrew, R. (2018). What is Justified Credence? *Episteme*.
<https://doi.org/10.1017/epi.2018.50>

Peer reviewed version

Link to published version (if available):
[10.1017/epi.2018.50](https://doi.org/10.1017/epi.2018.50)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Cambridge University Press at <https://www.cambridge.org/core/journals/episteme/article/what-is-justified-credence/0E5AA2D0F5C176C9BFCB3731CB0E0124> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

What is justified credence?

July 24, 2018

Abstract

In this paper, we seek a reliabilist account of justified credence. Reliabilism about justified beliefs comes in two varieties: process reliabilism (Goldman, 1979, 2008) and indicator reliabilism (Alston, 1988, 2005). Existing accounts of reliabilism about justified credence comes in the same two varieties: Jeff Dunn (2015) proposes a version of process reliabilism, while Weng Hong Tang (2016) offers a version of indicator reliabilism. As we will see, both face the same objection. If they are right about what justification is, it is mysterious why we care about justification, for neither of the accounts explains how justification is connected to anything of epistemic value. We will call this the Connection Problem. I begin by describing Dunn's process reliabilism and Tang's indicator reliabilism. I argue that, understood correctly, they are, in fact, extensionally equivalent. That is, Dunn and Tang reach the top of the same mountain, albeit by different routes. However, I argue that both face the Connection Problem. In response, I offer my own version of reliabilism, which is both process and indicator, and I argue that it solves that problem. Furthermore, I show that it is also extensionally equivalent to Dunn's reliabilism and Tang's. Thus, I reach the top of the same mountain as well.

Aafira and Halim are both 90% confident that it will be sunny tomorrow. Aafira bases her credence on her observation of the weather today and her past experience of the weather on days that follow days like today—around nine out of ten of them have been sunny. Halim bases his credence on wishful thinking—he just really likes the sun. Aafira, it seems, is justified in her credence, while Halim is not. Just as one of your full beliefs might be justified if it is based on visual perception under good conditions, or on memories of recent important events to which you attended carefully, or on testimony from experts you know to be

reliable, so might one of your credences be; and just as one of your full beliefs might be unjustified if it is based on wishful thinking, or on racially-biased stereotypical associations, or on testimony from ideologically driven news outlets, so might one of your credences be. In this paper, we seek an account of justified credence—in particular, we seek necessary and sufficient conditions for a credence to be justified. Our account will be reliabilist.

Reliabilism about justified beliefs comes in two varieties: process reliabilism (Goldman, 1979, 2008) and indicator reliabilism (Alston, 1988, 2005). Roughly, process reliabilism says that a belief is justified if it is formed by a reliable process, while indicator reliabilism says that a belief is justified if it is based on a ground that renders it likely to be true. Reliabilism about justified credence also comes in two varieties; indeed, it comes in the same two varieties. And, in fact, of the two existing proposals, Jeff Dunn's is a version of process reliabilism (Dunn, 2015) while Weng Hong Tang offers a version of indicator reliabilism (Tang, 2016). As we will see, both face the same objection. If they are right about what justification is, it is mysterious why we care about having justified doxastic attitudes, for neither of the accounts explains how justification is connected to anything of epistemic value. We will call this the *Connection Problem*. It begins from the observation that we value justified belief more than we value unjustified belief. And it proceeds by noting that, if Dunn's account of justified credence is correct, or if Tang's is, then this fact calls out for an explanation that neither of those accounts can provide.

I begin by describing Tang's indicator reliabilism and Dunn's process reliabilism. I argue that, once fixed up correctly in the light of certain issues, they are, in fact, extensionally equivalent. That is, Tang and Dunn reach the top of the same mountain, albeit by different routes. However, I argue that both face the Connection Problem. In response, I offer my own version of reliabilism, which is both process and indicator, and I argue that it solves that problem. Furthermore, I show that it is also extensionally equivalent to Tang's reliabilism and Dunn's. Thus, I reach the top of the same mountain as well.

1 Tang's indicator reliabilism for justified credence

Let's begin with Tang's indicator reliabilism for justified credences. According to indicator reliabilism for justified belief, a belief is justified if the ground on which it is based is a good indicator of the truth of that belief. Thus, beliefs formed on the basis of visual experiences tend to be justified because the fact that the agent had the visual experience in question makes it likely that the belief they based on it is true.¹ Wishful thinking, on the other hand, usually does not give rise to justified belief because the fact that an agent hopes that a particular proposition will be true—which in this case is the ground of their belief—does not make it likely that the proposition is true.² Tang follows William Alston's account of grounds, and I will too: the ground for a belief "is not what we might call the total complete input to the belief-forming mechanism, but rather those features of that input that are actually taken account of in forming the belief" (Alston, 1988, 268).

Tang seeks to extend this account of justified belief to the case of credence. Here is his first attempt at an account (Tang, 2016, 81).

Tang's Indicator Reliabilism for Justified Credence (first attempt)

A credence of x in X by an agent S is justified iff

(TIC1- α) it is actually based on ground g ;

(TIC2- α) the objective probability of X given that the agent has ground g approximates or equals x —we write this $P(X|S \text{ has } g) \approx x$.

Thus, just as an agent's full belief in a proposition is justified if its ground makes the objective probability of that proposition close to 1, a credence x in a proposition is justified if its ground makes the objective probability of that proposition close to x . There is a substantial problem here in identifying the notion of objective probability to which Tang wishes to appeal. But we will not address that here, other than to say that he conceives of it along the lines of

¹Obvious exceptions: when the lighting conditions are poor; when the agent is under the influence of hallucinogenics.

²Obvious exceptions: when, by hoping that you will win a competition, you make your success more likely, because your hope makes you perform better.

hypothetical frequentism—that is, the objective probability of X given Y is the hypothetical frequency with which propositions like X are true when propositions like Y are true.

Tang notes that, as it is stated, his version of indicator reliabilism faces a problem. Suppose I am presented with an urn. I know that it is filled with 100 balls, numbered 1 to 100, half of which are white, and half of which are black; but I do not know which are black and which white. I shake the urn vigorously and extract a ball. It's number 73 and it's white. I look at its colour and the number printed on it. I have a visual experience of a white ball with the numeral '73' on it. On the basis of my visual experience of the numeral alone, I assign credence 0.5 to the proposition that ball 73 is white. According to Tang's first version of indicator reliabilism for justified credence, my credence is justified. My ground is the visual experience of the number on the ball; I have that ground; I base my credence on that ground; and the objective probability that ball 73 is white given that I have a visual experience of the numeral '73' printed on it is 50%—after all, half the balls are white. Of course, the problem is that I have not used my total evidence—or, in the language of grounds, I have not based my belief on my most inclusive ground. I had the visual experience of the numeral on the ball as a ground; but I also had the visual experience of the numeral on the ball *and the colour of the ball* as a ground. The resulting credence is unjustified because the objective probability that ball 73 is white given I have the more inclusive ground is not 0.5—it is close to 1, since my visual system is very reliable. This leads Tang to amend his account of justified credence as follows, where we write $g' \subseteq g$ when g and g' are grounds, and g' is at least as inclusive as g —that is, g' is at least as strong as g , so that if g is the ground *visual experience of '73' printed on a ball* and g' is the ground *visual experience of numeral '73' printed on a white ball*, then $g' \subseteq g$; indeed, in this case, $g' \subsetneq g$, since g' is strictly stronger or more inclusive than g (Tang, 2016, 88).

Tang's Indicator Reliabilism for Justified Credence

A credence of x in X by an agent S is justified iff

(TIC1) it is actually based on ground g ;

(TIC2) if S has ground $g' \subseteq g$, then the objective probability of X given that the

agent has ground g' approximates or equals x —that is, $P(X|S \text{ has } g') \approx x$.

That is, a credence of x in X by S based on g is justified if the objective probability of X given that the agent has g is approximately x , and remains so as we condition on S having any more inclusive grounds g' that she also has. This, then, is Tang's version of indicator reliabilism for justified credences.

However, this, too, faces a problem, though this time one that Tang does not address: it distinguishes between cases that intuitively seem to warrant the same verdict. Suppose, for instance, that a die is rolled. Our friend, Persi, looks at it to see the number on which it landed. He tells me the following three pieces of information: (i) The die is fair; (ii) the number is less than 4; (iii) the number is greater than 1. He tells you this: (a) The die is fair; (b) the number is either 2 or 3. Thus, I have the following grounds:

$g_0 = \textit{The die is fair};$

$g_1 = \textit{The die is fair and the number is less than 4—that is, it is 1, 2, or 3};$

$g_2 = \textit{The die is fair and the number is greater than 1—that is, it is 2, 3, 4, or 5};$

$g_3 = \textit{The die is fair and the number is between 1 and 4—that is, it is 2 or 3}.$

So $g_3 \subseteq g_2 \subseteq g_0$ and $g_3 \subseteq g_1 \subseteq g_0$. On the other hand, your grounds are just g_0 and g_3 , where $g_3 \subseteq g_0$. You and I both have credence 0.5 in the proposition *Even*, which says that the die landed on an even number. We both base that credence on ground g_0 . Now, according to Tang's revised version of indicator reliabilism for justified credence, your credence is justified, while mine is not. After all, conditional on either of your grounds—the least inclusive g_0 and the most inclusive g_3 —the probability of *Even* is 0.5; whereas conditional on one of my two further grounds—the middle grounds g_1 —the probability of *Even* is 0.333. But that seems wrong. You might think that both are justified because the probability of *Even* conditional on the ground on which the credence is *in fact based*—namely, g_0 —matches the credence. Or you might think both are unjustified because, while the credence matches the probability of *Even* conditional on the most inclusive ground—namely, g_3 —which encodes the agent's total evidence, that is no credit to the agent, who based their belief on the less

inclusive ground g_0 —they just got lucky, you might say. But I can see no reason to think that one is justified and the other is not.

So I think Tang’s response to the example of the white ball marked ‘73’ is misguided. Rather, that example reveals that there are really two notions of justification. I will call them *actual-grounds justification* and *strongest grounds justification*. In the example, my credence of 0.5 that the ball is white is justified in the first sense, since the actual ground on which it is based indicates its truth to exactly the degree to which it is believed. But that credence is not justified in the second sense, since the strongest or most inclusive ground I have indicates its truth to a much higher degree than the degree to which it is believed. Thus, we have the following two versions of indicator reliabilism for justified credence:

Indicator Reliabilism for Actual-Grounds Justified Credence

A credence of x in X by an agent S is actual-grounds justified iff

(TIC1_{ag}) it is actually based on ground g ;

(TIC2_{ag}) $P(X|S \text{ has } g) \approx x$.

Indicator Reliabilism for Strongest-Grounds Justified Credence

A credence of x in X by an agent S is strongest-grounds justified iff

(TIC1_{sg}) g is the most inclusive ground that S has;

(TIC2_{sg}) $P(X|S \text{ has } g) \approx x$.

Of course, given the phenomenon of defeat, perhaps we should not count actual-grounds justification as a species of justification at all. I leave that discussion for another time. Here, I simply spell out that notion of justification in its indicator and process reliabilist forms, alongside the corresponding forms of strongest-grounds justification.

2 Reliabilism and Dunn on reliable credence

Next, let us turn to Dunn’s process reliabilism for justified credences. Now, to be clear, Dunn takes himself only to be providing an account of reliability for credences and credence-

forming processes. He doesn't endorse the further claim that a credence is justified just in case it is reliable. Instead, Dunn speculates that perhaps reliability is just one from an array of epistemic virtues, each of which is required for justification. Nonetheless, I will consider a version of reliabilism for justified credences that is based on Dunn's account of reliable credence.

For Dunn, a credence-forming process is perfectly reliable if it is well calibrated (or nearly so). Here's what it means for a process ρ to be well calibrated (Dunn, 2015, 1941-2).

- First, we construct a set of all and only the outputs of the process ρ in the actual world and in nearby counterfactual scenarios. An output of ρ consists of a credence x in a proposition X at a particular time t in a particular possible world w —so we represent it by the tuple (x, X, w, t) . If w is a nearby world and t a nearby time, we call (x, X, w, t) a nearby output. Let O_ρ be the set of nearby outputs—that is, the set of tuples (x, X, w, t) , where w is a nearby world, t is a nearby time, and ρ assigns credence x to proposition X in world w at time t .
- Second, we say that the truth-ratio of ρ for credence x is the proportion of nearby outputs (x, X, w, t) in O_ρ such that X is true at w and t .³
- Finally, we say that ρ is well calibrated (or nearly so) if, for each credence x that ρ assigns, x is equal to (or approximately equal to) the truth-ratio of ρ for x .⁴

For instance, suppose a process only ever assigns credence 0.6 or 0.7. And suppose that, 60% of the time that it assigns 0.6 in the actual world or a nearby world it assigns it to a proposition that is true; and 70% of the time it assigns 0.7 it assigns it to a true proposition. Then that process is well calibrated. If, on the other hand, it assigns 0.6 to a true proposition only 59% of the time and 0.7 to a true proposition 71% of the time, it is not well calibrated,

³In symbols:

$$TR(\rho, x) = \frac{|\{(x, W, w, t) \in O_\rho : X \text{ is true at } w \text{ and } t\}|}{|\{(x, X, w, t) \in O_\rho\}|}$$

⁴That is:

- ρ is well calibrated if $TR(\rho, x) = x$ for all x ;
- ρ is nearly well calibrated if $TR(\rho, x) \approx x$ for all x .

but it is nearly so. And if it assigns 0.6 to a truth 95% of the time and 0.7 to a truth only 5% of the time, then it is neither well-calibrated nor nearly so.

This, then, is Dunn's calibrationist account of the reliability of a credence-forming process—a credence-forming process is reliable iff it is well calibrated or nearly so. Any version of reliabilism about justified credences that is based on it requires two further ingredients. First, we must use the account to say when an individual credence is reliable; second, we must add the claim that a credence is justified iff it is reliable. Each of these creates problems.

Given Dunn's account of reliable credence-forming processes, when is an individual credence reliable? The natural answer: when it is formed by a reliable credence-forming process. But then we must be able to identify, for a given credence, the credence-forming process of which it is an output. The problem is that, for any credence, there are a great many processes of which it might be the output. I have a visual experience of a piece of red cloth on my desk, and I form a high credence that there is a piece of red cloth on my desk. Is this credence the output of a process that assigns a high credence that there is a piece of red cloth on my desk *whenever I have that visual experience*? Or is it the output of a process that assigns a high credence that there is a piece of red cloth on my desk *whenever I have that visual experience and the lighting conditions in my office are good*, while it assigns a middling credence that there is a piece of red cloth on my desk *whenever I have that visual experience and the lighting conditions in my office are bad*? It is easy to see that this is important. The first process is poorly calibrated, and thus unreliable on Dunn's account; the second process is better calibrated and thus more reliable on Dunn's account. This is the so-called *Generality Problem*, and it is a challenge that faces any version of reliabilism (Goldman, 1979; Feldman, 1985). I will offer a version of Juan Comesaña's solution to this problem (Comesaña, 2006).

Every account of justification must recognize that, for some important sense of justification, which we have been calling *actual grounds justification*, it is possible that two agents have the same doxastic attitude and the same evidence, while the doxastic attitude of one is justified and the doxastic attitude of the other is not, because they do not base that doxastic attitude on the same evidence. One agent might base their belief on the total evidence, for instance, whilst the other ignores that evidence and bases their belief purely on wishful

thinking. Thus, Comesaña claims, every theory of justification needs a notion of the *grounds* or *basis* of a doxastic attitude. But, once we have that, a solution to the Generality Problem is very close. Comesaña spells out the solution for process reliabilism about full beliefs:

Well-Founded Process Reliabilism for Actual-Grounds Justified Full Beliefs

A belief that X by an agent S is actual-grounds justified iff

(WPB1_{ag}) it is actually based on ground g ;

(WPB2_{ag}) the process *producing a belief state X based on ground g* is a reliable process.

This is easily adapted to the credal case:

Well-Founded Process Reliabilism for Actual-Grounds Justified Credences

A credence of x in X by an agent S is actual-grounds justified iff

(WPC1_{ag}) it is actually based on ground g ;

(WPC2_{ag}) the process *producing a credence of x in X based on ground g* is a reliable process.

Let us now try to apply Comesaña's solution to the Generality Problem to help Dunn's calibrationist reliabilism about justified credences. Recall: according to Dunn, a process ρ is reliable if it is well calibrated, or nearly so. Consider the process *producing a credence of x in X based on ground g* —for convenience, we'll write it $\rho_{X,x}^g$. There is only one credence that it assigns, namely x . So it is well calibrated if that truth-ratio of $\rho_{X,x}^g$ for x is equal to x . Now, $O_{\rho_{p,x}^g}$ is the set of tuples (X, x, w, t) where w is a nearby world and t a nearby time where $\rho_{X,x}^g$ assigns credence x to proposition X . But, by the definition of $\rho_{X,x}^g$, those are the nearby worlds and nearby times at which the agent has the ground g . Thus, the truth-ratio of $\rho_{X,x}^g$ for x is the proportion of those nearby worlds and times at which the agent has the ground g at which X is true. And that, it seems to me, is something like the objective probability of X conditional on the agent having ground g , at least given the hypothetical frequentist account of objective probability of the sort that Tang favours. As above, we denote the objective probability of X conditional on the agent S having grounds g as follows: $P(X|S \text{ has } g)$.

Thus, $P(X|S \text{ has } g)$ is the truth-ratio of $\rho_{p,x}^g$ for x . And thus, a credence x in X based on ground g is reliable on the calibrationist account iff x is close to $P(X|S \text{ has } g)$. That is,

Well-founded Calibrationist Process Reliabilism for Actual-Grounds Justified Credences

A credence of x in X by an agent S is actual-grounds justified iff

(WCPC1_{ag}) it is actually based on ground g ;

(WCPC2_{ag}) the process *producing a credence of x in X based on ground g* is a reliable process—that is, $P(X|S \text{ has } g) \approx x$.

But now compare this version of process reliabilism for actual-grounds justified credence, based on Dunn’s account of reliable processes and Comesaña’s solution to the Generality Problem, with Tang’s account of indicator reliabilism for actual-grounds justified credence. Consider the necessary and sufficient conditions that each imposes for justification: $\text{TIC1}_{ag} = \text{WCPC1}_{ag}$ and $\text{TIC2}_{ag} = \text{WCPC2}_{ag}$. Thus, these are the same account.

What’s more, we can easily formulate a version of process reliabilism for strongest-grounds justified credence that is also based on Dunn’s account of reliable processes and Comesaña’s solution to the Generality Problem.

Well-Founded Calibrationist Process Reliabilism for Strongest-Grounds Justified Credences

A credence of x in X by an agent S is strongest-grounds justified iff

(WCPC1_{sg}) g is the most inclusive ground that S has;

(WCPC2_{sg}) the process *producing a credence of x in X based on ground g* is a reliable process—that is, $P(X|S \text{ has } g) \approx x$.

Since TIC3_{sg} is equivalent to WCPC3_{sg} , this final version of process reliabilism is equivalent to Tang’s version of indicator reliabilism for strongest-grounds justified credences. Thus, Dunn and Tang have reached the top of the same mountain, albeit by different routes.

3 Two other routes up the mountain

Thus, we have now seen Dunn's process reliabilism and Tang's indicator reliabilism for both actual-grounds and strongest-grounds justified credences. Is either correct? If so, which? In one sense, both are correct; in another, neither is. Less mysteriously: Dunn's process reliabilism and Tang's indicator reliabilism are extensionally equivalent—that is, the same credences are justified on both accounts. This gives us a little hope that both have hit upon the correct account of justification. And indeed both are extensionally equivalent to the correct account of justified credence, which is thus a version of both process and indicator reliabilism. However, while they get the extension right, they do so for the wrong reasons. It may well be that a credence is justified just in case it is formed by a well calibrated process; but it is not justified *because* it is formed by a well calibrated process. And it may well be that a credence is justified just in case it matches the objective chance given its grounds; but it is not justified *because* it has that feature. So both have indeed hit upon the correct *extension* of the concept of justified credence; but they have done so for the wrong reasons, for they have not hit upon the correct *intension*. In the remainder, I will offer what I take to be the correct intension.

There are three sorts of route you might take when pursuing an account of justification for a given sort of doxastic attitude, such as a credence or a full belief. One: you might look to intuitions concerning particular cases and try to discern a set of necessary and sufficient conditions that sort these cases in the same way that your intuitions do. Two: you might begin with an existing account of justification for another sort of doxastic attitude and formulate your account for the sort of attitude that interests you by analogy. Three: you might begin with an account of epistemic value, assume that justification must be linked in some natural way to the promotion of epistemic value, and then provide an account of justification that vindicates that assumption. Dunn and Tang have both taken routes of the second sort, though both also make some appeal to intuitions. They take themselves to be generalising reliabilist accounts of justified full belief. In Tang's case, Alston's; in Dunn's case, Goldman's. I will follow a route of the third sort.

I will adopt the veritist's account of epistemic value. In the case of categorical doxastic attitudes—that is, belief, disbelief, suspension of judgment—veritism says that, if such an attitude is directed towards a true proposition, it is most valuable if it is a belief, least valuable if it is a disbelief, and neutrally valuable if it is a suspension of judgment; if the attitude is directed towards a false proposition, the order is reversed. In the case of graded doxastic states—that is, credences—veritism says that, if such an attitude is directed towards a truth, it is better the higher it is; and if it is directed towards a falsehood, it is better the lower it is. Given this account of epistemic value, what is the natural account of justification? Well, at first sight, there are two: one is process reliabilist; the other is indicator reliabilist. But, in a twist that should come as little surprise given the conclusions of the previous section, it will turn out that these two accounts coincide, and indeed coincide with the final versions of Dunn's and Tang's accounts that we reached above. Thus, I too will reach the top of the same mountain, but by another pair of routes.

3.1 Epistemic value version of indicator reliabilism

In the case of full beliefs, indicator reliabilism says this: a belief in X by S on the basis of grounds g is justified iff the objective probability of X given that S has grounds g is high—that is, close to 1. Tang's indicator reliabilism for justified credence is inspired by this. It is an attempt to generalise this account to the case of credence. However, I think he generalises in the wrong direction; that is, he takes the wrong feature to be salient and uses that to formulate his indicator reliabilism for justified credence. He takes the general form of indicator reliabilism to be something like this: a doxastic attitude s towards X by S on the basis of grounds g is justified iff the attitude s 'matches' (or comes close to 'matching') the objective probability of X given that S has grounds g . And he takes belief to 'match' high objective probability, and credence x to 'match' objective probability of x . The problem with this account is that it leaves mysterious why justification is valuable. Unless we say that matching objective probabilities is somehow epistemically valuable in itself, it isn't

clear why we should want to have justified doxastic attitudes in this sense.⁵ After all, there are many quantities pertaining to X and g that a credence might match. It might match the proportion of people with ground g who form a credence in X . Or it might match the proportion of people with credences in X who formed them on the basis of ground g . And so on. What is epistemically relevant about matching the objective probability of X given that you have ground g that isn't relevant about matching those other quantities? Call this *the connection problem*.

I contend instead that the general form of indicator reliabilism is this. (We will focus here on actual-grounds justification, but it will be clear how to adapt what we say to strongest-grounds justification.)

Indicator Reliabilism for Actual-Grounds Justified Doxastic Attitude (epistemic value version) A doxastic attitude s towards proposition X by agent S is actual-grounds justified iff

(EIA1_{ag}) it is actually based on g ;

(EIA2_{ag}) for every doxastic attitude s' of the same sort as s , the objective expected epistemic value of attitude s' towards X given that S has g is at most (or not much more than) the objective expected epistemic value of attitude s towards X given that S has g .

Here, when we say that attitude s' is of the same sort as attitude s , we mean that both are credences in X , or both are categorical doxastic attitudes towards X . Thus, a full belief in X is of the same sort as a full disbelief in X or a suspension in X ; and a credence of 0.6 in X is of the same sort as a credence of 0.3 in X ; but a credence of 0.3 or 0.7 in X is of a different sort from a belief or a disbelief in X . And when we talk of the objective expected epistemic value of an attitude, we mean its expected epistemic value calculated using the objective probability function. Thus, for the veritist, attitude s towards X by S is actual-grounds justified if s is based on a ground g that S has, and s is the attitude towards X

⁵See (Hájek, ms) and (Pettigrew, 2012) for the claim that credences aim to match the objective probability just as belief aims to match the truth. On this account, credences have greater epistemic value the closer they come to achieving that aim. See (Pettigrew, 2016a, Section 9.4) for what I take to be a decisive refutation.

that has (nearly) the highest objective expected accuracy conditional on S having g . The parenthetical qualifications simply make the account slightly less demanding than it would otherwise be.

Let's consider this in the full belief case. We have:

Indicator Reliabilism for Actual-Grounds Justified Belief (epistemic value version)

A belief in proposition X by agent S is actual-grounds justified iff

(EIB1_{ag}) it is actually based on g ;

(EIB2_{ag}) (a) the objective expected epistemic value of *disbelief* in X , given that S has g , is at most (or not much more than) the objective expected epistemic value of *belief* in X , given that S has g ;

(b) the objective expected epistemic value of *suspension* in X , given that S has g , is at most (or not much more than) the objective expected epistemic value of *belief* in X , given that S has g .

To complete this, we need only provide the veritist's account of the epistemic value of these categorical doxastic attitudes. It is natural to say that a belief in a truth and disbelief in a falsehood have the same high epistemic value—following Kenny Easwaran, we denote this R (for 'getting it Right'), and assume $R > 0$ (Easwaran, 2016). And it is natural to say that a disbelief in a truth and belief in a falsehood have the same low epistemic value—again following Easwaran, we denote this $-W$ (for 'getting it Wrong'), and assume $W > 0$. And finally it is natural to say that suspension of belief in a truth has the same epistemic value as suspension of belief in a falsehood, and both have epistemic value 0. Following Easwaran, we assume that $W > R$ —that is, we disvalue getting things wrong more than we value getting things right.⁶ Now, suppose proposition X has objective probability p . Then the expected epistemic utility of different categorical doxastic attitudes towards X is given

⁶Easwaran notes that, if $W = R$, and if we calculate the epistemic value of a set of doxastic attitudes by summing the value of the individual attitudes, then it is just as valuable to believe both a proposition and its negation as it is to suspend judgment on both; and if $W < R$, it is better to believe both a proposition and its negation than to suspend judgement on both. Thus, he assumes $W > R$.

below:

- Expected epistemic value of belief in $X = p \cdot R + (1 - p) \cdot (-W)$.
- Expected epistemic value of suspension in $X = p \cdot 0 + (1 - p) \cdot 0$.
- Expected epistemic value of disbelief in $X = p \cdot (-W) + (1 - p) \cdot R$.

Thus, belief in X has greatest epistemic value amongst the possible categorical doxastic attitudes to X if $p > \frac{W}{R+W}$; disbelief in X has greatest epistemic value if $p < \frac{R}{R+W}$; and suspension in X has greatest value if $\frac{R}{R+W} < p < \frac{W}{R+W}$. At $p = \frac{W}{R+W}$, belief ties with suspension, while at $p = \frac{R}{R+W}$, disbelief ties with suspension (Easwaran, 2016; Dorst, ta). With this in hand, we have the following version of indicator reliabilism for justified beliefs:

Indicator Reliabilism for Actual-Grounds Justified Belief (veritist version)

A belief in X by agent S is actual-grounds justified iff

(EIB1_{ag}^{*}) it is actually based on g ;

(EIB2_{ag}^{*}) the objective probability of X given that S has g is (nearly) greater than $\frac{W}{R+W}$.

And of course this is simply a more explicit version of the standard version of indicator reliabilism. It is more explicit because it gives a particular threshold above which the objective probability of X given that S has g counts as ‘high’, and above which (or not much below which) the belief in X by S counts as justified. That threshold is $\frac{W}{R+W}$, and it is determined by how much more you disvalue getting things wrong than you value getting things right.

Note that this also gives a straightforward account of when a suspension of judgment is justified. Replace (EIB2_{ag}^{*}) with:

(EIS2_{ag}^{*}) the objective probability of X given that S has g is (nearly) between $\frac{W}{R+W}$ and $\frac{R}{R+W}$.

And it says when a disbelief is justified. Replace (EIB3) with:

(EID2_{ag}^{*}) the objective probability of X given that S has g is (nearly) less than $\frac{R}{R+W}$.

Next, let’s turn to indicator reliabilism for justified credence. Here’s the epistemic value version:

Indicator Reliabilism for Actual-Grounds Justified Credence (epistemic value version)

A credence of x in proposition X by agent S is actual-grounds justified iff

(EIC1_{ag}) it is actually based on g ;

(EIC2_{ag}) for every credence x' , the objective expected epistemic value of credence x' in X given that S has g is at most (or not much more than) the objective expected epistemic value of credence x in X given that S has g .

Again, to complete this, we need an account of epistemic value for credences. As noted above, the veritist holds that the epistemic value of a credence is given by its accuracy. There is a lot to be said about different potential measures of the accuracy of a credence. Such measures are often called *scoring rules* or *local accuracy measures*, and there are many that are used for a variety of different purposes in epistemology, statistics, and psychology (Savage, 1971; de Finetti, 1974; Joyce, 1998, 2009; Predd et al., 2009; Pettigrew, 2016a).⁷ Here, I will say only this: we assume that those measures are *continuous* and *strictly proper*. That is: first, the accuracy of a credence in a given proposition varies continuously with that credence; second, any probability p in a proposition X expects credence p in X to be more accurate than it expects any other credence $x \neq p$ in X to be.⁸ These assumptions are widespread in the literature on accuracy-first epistemology, and many of the central arguments in that area rest upon them (Greaves & Wallace, 2006; Predd et al., 2009; Moss, 2011; Pettigrew, 2013; Horowitz, 2014; Schoenfield, 2015; Pettigrew, 2016b; Levinstein, 2015; Pettigrew, ta). Given both assumptions, (EIC2_{ag}) is provably equivalent to:

(EIC2_{ag}^{*}) the objective probability of X given that the agent has ground g approximates or equals x —that is, $P(X|S \text{ has } g) \approx x$.

⁷As we will use it here, a scoring rule is a measure of the accuracy of a credence. It is a function $s : \{0,1\} \times [0,1] \rightarrow [-\infty,0]$, so that $s(1,x)$ is the accuracy of having credence x in a true proposition while $s(0,x)$ is the accuracy of having credence x in a false proposition. Sometimes in the literature, scoring rules are taken to be measures of *inaccuracy*; that is, they are the negatives of the functions considered here. Here is an example: $q(1,x) = -(1-x)^2$ and $q(0,x) = -x^2$. This is sometimes called the *quadratic scoring rule*.

⁸More precisely: $s(1,x)$ and $s(0,x)$ are continuous functions of x ; and, for any $0 \leq p \leq 1$, $ps(1,x) + (1-p)s(0,x)$ is maximised, as a function of x , at $p = x$. Recall the quadratic scoring rule q from footnote 7. It is clearly continuous, and a little calculus shows that $pq(1,x) + (1-p)q(0,x) = -p(1-x) - (1-p)x^2$ is maximised, as a function of x , at $p = x$.

But of course $\text{EIC2}_{ag}^* = \text{TIC2}_{ag} = \text{WCPC2}_{ag}$ from above. Thus, the veritist version of indicator reliabilism for actual-grounds justified credences is equivalent to Tang's indicator reliabilism, and to the calibrationist version of process reliabilism.

3.2 Epistemic value version of process reliabilism

Next, let's turn to process reliabilism. How might we give an epistemic value version of that? The mistake made by the calibrationist's version of process reliabilism is of the same sort as the mistake made by Tang in his formulation of indicator reliabilism—both generalise from the case of full beliefs in the wrong way by mistaking an accidental feature for the salient feature. For the calibrationist, a full belief is justified if it is formed by a reliable process, and a process is reliable if a high proportion of the beliefs it produces are true. Now, notice that there is a sense in which such a process is calibrated: a belief is associated with a high degree of confidence, and that matches, at least approximately, the high truth-ratio of a reliable process. In fact, we want to say that this process is *belief*-reliable. For it is possible for a process to be reliable in its formation of beliefs, but not in its formation of disbeliefs. A process is *disbelief*-reliable if a high proportion of the disbeliefs it produces are false. And we might say that a process is *suspension*-reliable if a middling proportion of the suspensions it forms are true and a middling proportion are false. In each case, we think that there is, corresponding to each sort of categorical doxastic attitude s , a fitting proportion x such that a process is s -reliable if x is (approximately) the proportion of truths amongst the propositions to which it assigns s . Applying this in the credal case gives us the calibrationist version of process reliabilism that we have already met—a credence x in S is justified if it is formed by a process whose truth-ratio for a given credence is equal to that credence.

However, being the product of a belief-reliable process is not the feature of a belief in virtue of which it is justified. After all, it is not a feature that is connected to epistemic value. Thus, the calibrationist account of justified belief or credence faces the same connection problem that we identified for Tang's account. Why should we want to have beliefs or credences that are justified in this sense? Why should we want to have beliefs or credences that are formed by well calibrated processes? What epistemic value does this promote?

One possibility is that forming beliefs or credences using a well calibrated process makes it more likely that your credences themselves will be well calibrated, something that some philosophers take to be a source of epistemic value (van Fraassen, 1983; Shimony, 1988; Lange, 1999). We say that a set of credences is well calibrated if, for each $0 \leq x \leq 1$, the proportion of true propositions to which we assign x is x . The problem is that, as is well known, being well calibrated is not in fact the aim of credence nor a source of epistemic value (Seidenfeld, 1985). For one thing, it is too easy to come by—providing you have a credence in the negation of a proposition whenever you have a credence in that proposition, you can ensure that your credences are well calibrated by assigning 0.5 to each proposition—whatever the truth values of the propositions, your credences will be well calibrated. Also, for an agent with credences only in a proposition and its negation, while credences of 0.5 in both are guaranteed to be well calibrated, a credence of 0.99 in the proposition and a credence of 0.01 in its negation are guaranteed not to be; thus, they would count as less epistemically valuable according to the calibrationist account of epistemic value, even if the proposition is true and its negation false.

Thus, being the product of a belief-reliable process is not the feature of a belief in virtue of which it is justified. Rather, I submit, a belief is justified if it is the product of a process that has high expected epistemic value.

Process Reliabilism for Actual-Grounds Justified Doxastic Attitude (epistemic value version)

Doxastic attitude s towards proposition X by agent S is actual-grounds justified
iff

(EPA1_{ag}) s is actually produced by a process ρ ;

(EPA2_{ag}) If ρ' is a process that is available to S , then the expected epistemic value of ρ' is at most the expected epistemic value of ρ .

To complete this account, we must say which processes count as available to an agent? To answer this, recall Comesaña's solution to the Generality Problem. On this solution, the only processes that interest us have the form, *process producing doxastic attitude s towards X on basis*

of ground g . Clearly, a process of this form is available to an agent exactly when the agent has ground g . This gives

Well-founded Process Reliabilism about Actual-Grounds Justified Doxastic Attitude (epistemic value version) Attitude s towards proposition X by S is actual-grounds justified iff

(WEPA1_{ag}) it is actually produced by process $\rho_{s,X}^g$;

(WEPA2_{ag}) for every doxastic attitude s' of the same sort as s , the expected epistemic value of process $\rho_{s',X}^g$ is at most the expected epistemic value of process $\rho_{s,X}^g$.

Thus, in the case of full beliefs, we have:

Process Reliabilism for Justified Belief (epistemic value version) A belief in proposition X by agent S is actual-grounds justified iff

(EPB1_{ag}) it is actually produced by process $\rho_{\text{bel},X}^g$;

(EPB2_{ag}) (a) the expected epistemic value of process $\rho_{\text{dis},X}^g$ is at most (or not much more than) the expected epistemic value of process $\rho_{\text{bel},X}^g$;

(b) the expected epistemic value of process $\rho_{\text{sus},X}^g$ is at most (or not much more than) the expected epistemic value of process $\rho_{\text{bel},X}^g$.

And it is easy to see that (EPB1_{ag}) = (EIB1_{ag}), since belief in X by S is produced by process $\rho_{\text{bel},X}^g$ iff S has ground g and her belief in X is based on g . Also, (EPB2_{ag}) is equivalent to (EIB2_{ag}). Thus, as for the epistemic version of indicator reliabilism, we get:

Process Reliabilism for Actual-Grounds Justified Beliefs (veritist version)

A belief in X by agent S is actual-grounds justified iff

(EPB1*_{ag}) it is actually based on g ;

(EPB2*_{ag}) the objective probability of X given that S has g is (nearly) greater than $\frac{W}{R+W}$.

Next, consider how the epistemic value version of process reliabilism applies to credences.

Process Reliabilism for Doxastically Justified Credence (epistemic value version)

A credence of x in proposition X by agent S is actual-grounds justified iff

(EPC1_{ag}) it is actually produced by process $\rho_{x,X}^g$;

(EPC2_{ag}) for any credence x' the expected epistemic value of process $\rho_{x',X}^g$ is at most (or not much more than) the expected epistemic value of process $\rho_{x,X}^g$.

As before, we see that (EPC1_{ag}) is equivalent to (EIC1_{ag}). And, providing the measure of accuracy is continuous and strictly proper, we get that (EPC2_{ag}) is equivalent to (EIC2_{ag}). So, once again, our route takes us to the same summit. The routes taken by Tang, Dunn, and the epistemic value versions of process and indicator reliabilism lead to the same spot, namely, the following account of actual-grounds justified credence:

Reliabilism for Actual-Grounds Justified Credence (epistemic value version)

A credence of x in proposition X by agent S is actual-grounds justified iff

(ERC1_{ag}) it is actually based on g ;

(ERC2_{ag}) the objective probability of X given that the agent has ground g approximates or equals x —that is, $P(X|S \text{ has } g) \approx x$.

And the following account of strongest-grounds justified credence:

Reliabilism for Strongest-Grounds Justified Credence (epistemic value version)

A credence of x in proposition X by agent S is strongest-grounds justified iff

(ERC1_{sg}) g is the most inclusive ground that S has;

(ERC2_{sg}) the objective probability of X given that the agent has ground g approximates or equals x —that is, $P(X|S \text{ has } g) \approx x$.

References

- Alston, W. (1988). An Internalist Externalism. *Synthese*, 74, 265–83.
- Alston, W. (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*. Ithaca, NY: Cornell University Press.
- Comesaña, J. (2006). A well-founded solution to the generality problem. *Philosophical Studies*, 129, 27–47.
- de Finetti, B. (1974). *Theory of Probability*, vol. I. New York: John Wiley & Sons.
- Dorst, K. (ta). Lockeans Maximize Expected Accuracy. *Mind*.
- Dunn, J. (2015). Reliability for degrees of belief. *Philosophical Studies*, 172(7), 1929–1952.
- Easwaran, K. (2016). Dr Truthlove, Or: How I Learned to Stop Worrying and Love Bayesian Probabilities. *Noûs*, 50(4), 816–853.
- Feldman, R. (1985). Reliability and Justification. *The Monist*, 68(2), 159–174.
- Goldman, A. (1979). What is justified belief? In G. S. Pappas (Ed.) *Justification and Knowledge*. Dordrecht: Reidel.
- Goldman, A. (2008). Immediate justification and process reliabilism. In Q. Smith (Ed.) *Epistemology: New Essays*, (pp. 63–82). New York: Oxford University Press.
- Greaves, H., & Wallace, D. (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459), 607–632.
- Hájek, A. (ms). A Puzzle about Partial Belief.
- Horowitz, S. (2014). Immoderately rational. *Philosophical Studies*, 167, 41–56.
- Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4), 575–603.
- Joyce, J. M. (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In F. Huber, & C. Schmidt-Petri (Eds.) *Degrees of Belief*. Springer.

- Lange, M. (1999). Calibration and the Epistemological Role of Bayesian Conditionalization. *The Journal of Philosophy*, 96(6), 294–324.
- Levinstein, B. A. (2015). With All Due Respect: The Macro-Epistemology of Disagreement. *Philosophers' Imprint*, 15(3), 1–20.
- Moss, S. (2011). Scoring Rules and Epistemic Compromise. *Mind*, 120(480), 1053–1069.
- Pettigrew, R. (2012). Accuracy, Chance, and the Principal Principle. *Philosophical Review*, 121(2), 241–275.
- Pettigrew, R. (2013). A New Epistemic Utility Argument for the Principal Principle. *Episteme*, 10(1), 19–35.
- Pettigrew, R. (2016a). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pettigrew, R. (2016b). Jamesian epistemology formalised: an explication of 'The Will to Believe'. *Episteme*, 13(3), 253–268.
- Pettigrew, R. (ta). On the Accuracy of Group Credences. In T. S. Gendler, & J. Hawthorne (Eds.) *Oxford Studies in Epistemology*, vol. 6. Oxford: Oxford University Press.
- Predd, J., Seiringer, R., Lieb, E. H., Osherson, D., Poor, V., & Kulkarni, S. (2009). Probabilistic Coherence and Proper Scoring Rules. *IEEE Transactions of Information Theory*, 55(10), 4786–4792.
- Savage, L. J. (1971). Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, 66(336), 783–801.
- Schoenfield, M. (2015). Bridging Rationality and Accuracy. *Journal of Philosophy*, 112(12), 633–657.
- Seidenfeld, T. (1985). Calibration, Coherence, and Scoring Rules. *Philosophy of Science*, 52(2), 274–294.
- Shimony, A. (1988). An Adamite Derivation of the Calculus of Probability. In J. Fetzer (Ed.) *Probability and Causality*. Reidel.

Tang, W. H. (2016). Reliability Theories of Justified Credence. *Mind*, 125(497), 63–94.

van Fraassen, B. C. (1983). Calibration: Frequency Justification for Personal Probability.
In R. S. Cohen, & L. Laudan (Eds.) *Physics, Philosophy, and Psychoanalysis*. Dordrecht:
Springer.