

Reliability Theories of Justified Credence

Weng Hong Tang

April 2014

Abstract

Reliabilists hold that a belief is doxastically justified if and only if it is caused by a reliable process. But since such a process is one that tends to produce a high ratio of true to false beliefs, reliabilism is on the face of it applicable to *binary beliefs*, but not to *degrees of confidence* or *credences*. For while (binary) beliefs admit of truth or falsity, the same cannot be said of credences in general. A natural question now arises: can reliability theories of justified belief be extended or modified to account for justified credence? In this paper, I address the preceding question. I begin by showing that, as it stands, reliabilism cannot account for justified credence. I then consider three ways in which the reliabilist may try to do so by extending or modifying her theory, but I argue that such attempts face certain problems. After that, I turn to a version of reliabilism that incorporates *evidentialist* elements and argue that it allows us to avoid the problems that the other theories face. If I am right, this gives reliabilists a reason, aside from those given recently by Comesaña (2010) and Goldman (2011), to move towards such a kind of hybrid theory.

1 Introduction

We are familiar with the notion of *binary* belief. Suppose we are considering the proposition that the train will be late. Either we believe the proposition or we do not. If the latter, then either we disbelieve the proposition or we suspend judgement in it. But there is also room for a more fine-grained notion of belief, namely, that of *degree of confidence* or *credence*. For example, we may be absolutely confident that two plus two equals four, have very little confidence that we will win the lottery, be more confident that two plus two equals four than that the train will be late, be thrice as confident that a fair coin will come up heads than that a fair die will roll a six, and so on. Given that confidence or credence comes in varying strengths, it is useful to represent its strength with a number between 0 and 1. Suppose we form an opinion about some proposition p . Our credence in p is 1 if we are absolutely confident that p , and 0 if we are absolutely confident that $\neg p$. But it may also equal some value bigger than 0 and smaller than 1. For example, we may assign a credence of 0.8 to the proposition that the train will be late—we may be rather confident, though not absolutely sure, that such a proposition is true. Or we may be as confident that a coin will land heads as we are that it will not do so, in which case our credence in the coin landing heads is 0.5.

Credences, like binary beliefs, may be justified or unjustified. For example, suppose I have overwhelming evidence which suggests that it is raining and that the streets will be wet if so. If I then form a low credence in the streets being wet, my low credence seems unjustified. But how should we account for the justifiedness of our credences? In this paper, I address the preceding question with respect to *process reliabilism*, a prominent theory of justification, as well as arguably the most popular and most discussed *reliability* theory of justification.

The plan is as follows. In the next section, I say a bit more about traditional process reliabilism. Though it tells a familiar story about the justifiedness of binary beliefs, I argue that it faces difficulty, at least as it stands, in accounting for justified

credence. In section 3, I discuss three proposals on how to extend or modify the theory to accommodate credences. But I argue that such proposals face problems. Then in section 4, I turn to a version of process reliabilism which is inspired by Alston (2005) and which is notably different from the other proposals. This account incorporates certain *evidentialist* elements that, I argue, help us avoid the problems that the other proposals face. If I am right, this gives reliabilists a reason, aside from those given recently by Comesaña (2010) and Goldman (2011), to move towards such a kind of hybrid theory. In section 5, I conclude.

2 Traditional process reliabilism and credences

On pain of circularity, when reliabilists spell out the conditions under which a belief is justified, they typically avoid invoking epistemic notions such as *knowledge*, *justification*, *epistemic probability*, *rationality*, and the like. Instead, they invoke non-epistemic notions such as *truth* or the *tendency* of a process to *cause* or *produce* true beliefs. According to process reliabilism, its most prominent advocate being Goldman (1979), a belief is doxastically justified if and only if it is produced—or sustained—by a reliable process, that is, one that tends to cause a high proportion of true to false beliefs.¹ For example, suppose that my belief that I will soon become a millionaire is caused by wishful thinking. Then it is unjustified because wishful thinking tends not to produce a high proportion of true beliefs. Or suppose that my belief that it is raining is caused by my looking out the window and carefully observing the environment. Then it is justified because perception and careful observation tend to lead to a high proportion of true beliefs.

Instead of formulating reliabilism in terms of a single process, the process reliabilist sometimes formulates it in terms of ‘an entire system of processes’ (Goldman 1988, pp. 63–4). This formulation takes into account cases in which a justified

¹A theory of *doxastic* justification attempts to spell out the conditions under which a belief that we have is justified, whereas a theory of *propositional* justification attempts to spell out the conditions under which there is justification for a belief, whether or not we have the belief. For more on the distinction between doxastic and propositional justification, see Firth 1978. Henceforth, unless otherwise specified, the term ‘justification’ should be understood to refer to doxastic justification.

belief is formed by a process that is not reliable per se, but is *conditionally reliable* (Goldman 1979). Consider, for example, the process of reasoning by *modus ponens*, which takes certain beliefs as input in order to form a new belief. Since whether the outputs of the process are true depends on whether the inputs are true, it is neither reliable nor unreliable per se. But it is conditionally reliable—it will tend to produce true beliefs given that the beliefs taken as input are true. Now, suppose that the input consists of a belief that if p then q and a belief that p , and the output is a belief that q . The latter belief may not be justified if either of the former two is not. But it may be justified if, say, reliable processes led to a belief that if p then q and a belief that p , and a conditionally reliable process then took these two beliefs as input to produce the belief that q .

Though traditional process reliabilism tells a familiar story with respect to the justifiedness of binary beliefs, it does not on the face of it accommodate credences well. For it does not make sense to say that a reliable process tends to produce a high proportion of true credences. Whereas beliefs admit of truth or falsity, the same cannot be said of credences in general—witness, for example, the oddity of saying that a credence of 0.5 in the streets being wet is true or false.

It would be too quick, however, to conclude on the basis of such an observation that reliabilists cannot account for justified credence. Granted, a credence of 0.5 is not something that is either true or false. But there still seems to be something right about having such a credence if, for instance, the corresponding objective probability is also 0.5. Might there not be a simple way for the reliabilist to account for justified credence by exploiting this idea? For example, might we not simply hold that a credence is justified if and only if it is produced by a process that tends to produce credences which match the corresponding objective probabilities?

Indeed, the above proposal is a natural one to make and deserves serious consideration. I consider it, as well as two other proposals on how to extend or modify reliabilism to account for justified credence, in the next section. But first, let me discuss briefly a suggestion that avoids any significant amendment to reliabilism.

Though credences are not in general true or false, their contents—the propositions to which we assign such credences—may be so. Given this, one may suggest that a reliable process is one that tends to lead to a high proportion of credences with true contents. But while this suggestion has the benefit of not requiring the reliabilist to amend her theory significantly, it suffers from being too indiscriminating. Suppose we know that the frequency of rain in our city is 90%. Despite that, for each day of the coming month, our credence-forming process assigns a credence of 0.51 to it raining that day. Suppose it will in fact rain every day next month. Then a high proportion of the credences will have true contents and—according to the suggestion—justified. But such credences are patently *not* justified. For any day in the coming month, our credence that it will rain that day should be way higher than 0.51.²

One may think that merely saying that our credence of 0.51 in rain is unjustified does not capture the epistemic situation described in the above example fully. For one may think that justification comes in degrees and that a credence of 0.51 in rain is still more justified than, say, a credence of 0.2 in rain.³ I take the point. But it remains that the current suggestion on how to account for justified credence is very unsatisfactory. First, it cannot account for why our credence of 0.51 in rain seems to be unjustified *tout court*. Second, granting that justification comes in degrees, a credence of 0.9 in rain seems to be more justified than a credence of 0.8 in rain, which in turn seems to be more justified than a credence of 0.51 in rain. The current suggestion cannot account for this too. (And I will argue in section 4 that my favoured reliability account of justified credence can do so.)

²The suggestion faces another worry, although there might be a way to deal with it. Consider a process that produces a credence of $1 - x$ in $\neg p$ whenever it produces a credence of x in p . Given that either p or $\neg p$, exactly half the credences produced by the process will have true contents, and exactly half will have false contents. So the process will *not* produce a high proportion of credences with true contents. Yet, the credences produced need not be unjustified for that. In response, one may hold that when determining the proportion of credences with true contents, we should just focus on credences with values greater than 0.5 in order to avoid double-counting. For one may think that a credence of x in p and a credence of $1 - x$ in $\neg p$ are really the same state of mind, or at least, that a *minimally rational* agent (in the sense used in Cherniak 1981) who assigns a credence of x to p will assign a credence of $1 - x$ to $\neg p$.

³I am grateful to an anonymous referee for raising this point.

3 Three reliability theories of justified credence

To account for justified credence, we cannot just rely on traditional process reliabilism—we need to extend or modify the theory. In this section, I consider three proposals on how to do so. I also raise some worries for them.

3.1 The theories: (Probability), (Calibration), and (Brier)

Even though credences are not in general true or false, one may think that there is an analogue of truth for credence. In particular, one may think that whereas belief aims at truth, credence aims at *objective probability* of some sort—that whereas a belief that p is *vindicated* if p is true, a credence of x in p is vindicated if the objective probability of p being true is x .⁴ Someone who holds such a view may like the following theory, briefly described in section 2:

(Probability) For any $x \in [0, 1]$ and any p , a credence of x in p is justified if and only if it is caused by a process that tends to result in a high proportion of credences that match or approximate the corresponding objective probabilities.

Invoking different kinds of objective probability—be it objective chance or propensity—will yield different versions of (Probability). Note that (Probability) does not require that, for the credences caused by a process to be justified, a high proportion of them must match the corresponding objective probabilities exactly. For requiring exact matches may be too demanding and mean that we hardly have any justified credences. Also, as formulated, (Probability) concerns whether a credence is justified *tout court*. But given that justification comes in degrees, we may hold that the higher (or lower) the proportion of credences that match or approximate the corresponding objective probabilities, the more (or less) justified such credences will be.

⁴Talk of vindication can be found in van Fraassen 1983, p. 296. It is also found in Hájek MS, where it is argued that whereas belief aims at truth, credence aims at objective chance.

Now, one may worry that the term ‘high’ is vague—how high must the relevant proportion of credences be in order for such credences to count as being justified? Relatedly, one may also wonder when a credence in p counts as approximating the corresponding objective probability of p . Does it count as doing so when their values differ by no more than 0.01, 0.05, or 0.1, etc.?⁵

In response to the first worry, let us grant that the term ‘high’ is vague. But such vagueness is not a problem. Lewis observes, in defending his theory of counterfactuals, that though the notion of comparative similarity is vague, it is precisely because of its vagueness that it is suitable for analysing counterfactuals, since the truth-conditions for counterfactuals admit of vagueness too (Lewis 1973, pp. 91–3). And Goldman writes:

Stated in its simplest form, process reliabilism says that a belief is justified if and only if it is caused by a sufficiently reliable process . . . , where *the threshold of sufficiency is left vague (as befits the notion of justification, which itself is vague)*. (Goldman 2012, p. 4, my emphasis)

In a similar vein, we may hold that appealing to the vague notion of a high proportion is useful precisely because whether a person counts as having a credence that is justified *tout court* is something that admits of vagueness too. Perhaps what counts as justified depends on context: the more stringent the epistemic standards are in a particular context, the higher the relevant proportion of credences needs to be for the credences to count as being justified.

In response to the second worry, let us grant that there may be no determinate answer to the question of when a credence counts as approximating the corresponding objective probability. But this is not a problem given that justification comes in degrees. Suppose a process produces a high proportion of credences that approximate the corresponding objective probabilities. The proponent of (Probability) may hold that the closer such credences approach the corresponding objective

⁵I am grateful to an anonymous referee for raising these questions.

probabilities, the more justified they are.⁶ And the degree of justification necessary for a credence to be considered justified *tout court* may vary with context. In a context with very stringent epistemic standards, a high proportion of credences may have to match the corresponding objective probabilities exactly for the credences to count as being justified. In another context, such credences may count as being justified only if they do not differ from the corresponding objective probabilities by more than a value of, say, 0.01, or 0.05, or 0.1 (depending on the stringency of epistemic standards).

We have seen one way in which the reliabilist may attempt to account for justified credence. However, instead of appealing to objective probabilities, one may appeal to the notion of *calibration*. To illustrate what it is for something to be *calibrated*, van Fraassen writes:

[C]onsider a weather forecaster who says in the morning that the probability of rain equals 0.8. That day it either rains or does not. How good a forecaster is he? Clearly to evaluate him we must look at his performance over a longer period of time. Calibration is a measure of agreement between judgments and actual frequencies . . . This forecaster was perfectly calibrated over the past year, for example, if, for every number r , the proportion of rainy days among those days on which he announced probability r for rain, equalled r . (van Fraassen 1984, p. 245)

Just as with weather forecasters, one may think that credence-forming processes should be evaluated by whether they are calibrated, where a process is calibrated (or approximately calibrated) just in case, for any y , $100y\%$ (or approximately $100y\%$) of the propositions in which the process causes us to have a credence of y are true. Appealing to the notion of a calibrated process, one may put forward the following theory:

⁶Cf. Goldman 2012, p. 26.

(Calibration) For any $x \in [0, 1]$ and any p , a credence of x in p is justified if and only if it is caused by a calibrated (or approximately calibrated) process.⁷

Since it is unlikely that our credence-forming processes will be perfectly calibrated, the qualification ‘approximately calibrated’ helps (Calibration) avoid the consequence that hardly any of our credences are justified.⁸ Admittedly, the term ‘approximately’ is vague. But this is not a problem, as seen from the earlier discussion of (Probability). Also, as formulated, (Calibration) deals only with credences that are justified or unjustified *tout court*. But given that justification comes in degrees, one may hold that better justified credences are produced by better calibrated processes.⁹

Instead of appealing to calibration or to objective probabilities, one may appeal to the notion of *accuracy*, where the greater the credence we assign to a true (or false) proposition, the more accurate (or inaccurate) it is. Following Joyce (1998), we may measure the accuracy of a credence of x in p using the *Brier score*, calculated with the formula $(x - T(p))^2$, where $T(p)$ equals 1 if p is true and 0 if p is false.¹⁰ In general, a lower Brier score corresponds to a more accurate credence. And at the extreme, a credence of x in p is perfectly accurate if $x = 0$ and p is false, or if $x = 1$ and p is true. For example, suppose that it will rain. Then assigning a credence of 0.8 to the proposition that it will rain will yield a Brier score of 0.04, while assigning a credence of 0.6 to the same proposition will yield a score of 0.16. The first credence, having a lower Brier score than the second, is more accurate than it. Also, a credence of 1 in the proposition that it will rain is perfectly accurate,

⁷Lam (2011, pp. 212–15) has recently considered a version of (Calibration). See also Goldman 2012, p. 26 and Goldman 1986, pp. 113–15. For more on the notion of calibration, see Lichtenstein et al. 1982.

⁸We may tweak (Calibration) to focus on *potential* rather than *actual* calibration (van Fraassen 1983, pp. 302–5). Suppose that van Fraassen’s weather forecaster is calibrated. Now consider a counterpart who makes his first and only prediction and then dies, but who would have continued making the same predictions as those made by the actual forecaster if he had not died. The counterpart is not in fact calibrated, but one might think that his first prediction is just as justified as that made by the actual forecaster. And one may support this view by noting that the counterpart is potentially calibrated—that if he had not died and had kept on making predictions similar to those made by the actual forecaster, he would have been calibrated.

⁹Cf. Goldman 2012, p. 26.

¹⁰Brier (1950) first proposed that we use the score to gauge the accuracy of weather forecasts.

whereas a credence of 0 in the same proposition is perfectly inaccurate.

Appealing to the notion of accurate credences, one may put forward the following theory:

(Brier) For any $x \in [0, 1]$ and any p , a credence of x in p is justified if and only if it is caused by a process that tends to result in credences that have a low average Brier score.¹¹

By cashing out reliability in terms of the ratio of true to false beliefs, traditional process reliabilism captures the idea that there is something good about believing truths and not believing falsehoods. By appealing to the notion of accuracy, (Brier) captures the analogous idea that there is something good about assigning high credences to truths and low credences to falsehoods. Note that to accommodate the view that justification comes in degrees, we may hold that the lower the average Brier score of the credences produced by a process, the more justified such credences will be.

3.2 Problems with (Probability), (Calibration), and (Brier)

We have seen three attempts to account for justified credence along reliabilist lines. The first appeals to a match between credences and the corresponding objective probabilities; the second appeals to the notion of calibration; and the third appeals to the notion of an accurate credence. But the attempts all face certain problems.

(Probability) faces the following problem. It seems possible that there exists

¹¹Lam (2011, pp. 215–219) has recently discussed a version of (Brier). Note that adopting different scoring rules will yield different theories that are nevertheless similar in flavour to (Brier). For example, Goldman (1999b, p. 90) proposes that we evaluate credences by their *veritistic values*, where for any x , the veritistic value of a credence of x in a true proposition equals x . Suppose we have a credence of 0.7 in p and a credence of 0.2 in q , where p is true and q is false. Then the veritistic values of our credences are 0.7 and 0.8 respectively (assuming that if we have a credence of x in a proposition, we also have a credence of $1 - x$ in its negation). Appealing to the notion of veritistic value, one may advance the following theory:

(Verity) For any $x \in [0, 1]$ and any p , a credence of x in p is justified if and only if it is caused by a process that tends to result in credences that have a high average veritistic value.

For reasons of economy, I will focus on (Brier)—what I say about it should apply *mutatis mutandis* to (Verity) or to any theory similar to (Brier) that invokes a similar scoring rule to assess the accuracy of credences.

an omniscient agent such as God who assigns credence 1 to *all* true propositions and credence 0 to *all* false ones and is perfectly justified in doing so. But (Probability) rules out such a possibility. For there are presumably many propositions whose objective probabilities of being true are *neither* high nor low, and given this, a significant proportion of the omniscient agent's credences will not approximate the relevant objective probabilities. By the lights of (Probability), the agent's doxastic states are not perfectly justified—an omniscient agent who assigns only extreme credences of 1 and 0 and is perfectly justified in doing so cannot exist. But even ardent atheists should be wary of such an easy argument against the existence of God. At any rate, it is best if the correctness of a theory of justified credence does not hinge on whether such an agent is possible.¹²

The case above involves an agent who can do better than merely assign credences that match the corresponding objective probabilities. And it seems that we should not be penalized for doing better. Next, let us consider another kind of case in which, based on the evidence available to an agent, she tends to assign credences that do not match the corresponding objective probabilities but is nevertheless justified in doing so. Suppose we initially have no evidence whatsoever as to the truth of a bunch of propositions and—unbeknownst to us—their objective probabilities of being true are either 1 or 0. Now, Jack comes along and asserts a number of the propositions. We know that Jack has a particular quirk—60% of the time, he will assert something he knows to be true, while 40% of the time, he will assert something he knows to be false. Unbeknownst to us, 60% of the propositions Jack

¹²One may hold that if future contingents—that is, contingent propositions about the future—are neither true nor false, then right now, no agent can know for certain whether such propositions are true or false. This is compatible with the existence of an omniscient agent in the sense defined above. For failing to assign a credence of 1 or 0 to a proposition that is neither true nor false is compatible with assigning credence 1 to all truths and credence 0 to all falsehoods. But then, as an anonymous referee points out, we have room for an omniscient God who assigns credences of intermediate values to some propositions. In response, let me grant that the above argument against (Probability) is not watertight. Nevertheless, it brings out what the proponent of (Probability) is committed to: she has to hold that future contingents are neither true nor false and subscribe to a view about God that goes against tradition. Given that such views are controversial, not all proponents of (Probability) may want to be saddled with them. (For more discussion of the view that even God does not know for certain whether some propositions are true or false, see Rissler 2006.)

has just asserted have an objective probability of 1 of being true, while the rest have an objective probability of 0 of being true. In such a case, based on what we know about Jack, assigning a credence of 0.6 to each proposition he has just asserted seems justified. But (Probability) says otherwise. For the credence-forming process responsible for such an assignment does not tend to produce credences that approximate the corresponding objective probabilities.

(Calibration) avoids the above problems. Whatever process is responsible for God's credences will be perfectly calibrated because all propositions to which God assigns a credence of 1 will be true, and all propositions to which God assigns a credence of 0 will be false. Also, a process that assigns a credence of 0.6 to each proposition that Jack has just asserted will be calibrated (at least with respect to those propositions).

But a different problem arises. As statisticians have pointed out, a perfectly calibrated weather forecaster need not always be the most informative.¹³ Suppose a forecaster knows that the long-run frequency of rain in a particular city is 0.2. And every day, she reports a 20% probability of rain in that city without looking at the latest and most relevant meteorological data. Despite being calibrated, her predictions will not be very useful to us. Similarly, calibration is not sufficient for justification. Suppose our credence-forming processes assign credences based simply on our knowledge of the corresponding long-run frequencies while ignoring other additional evidence that might come our way. (Calibration) will deem the credences produced by such processes justified since the processes will be approximately calibrated. But this ignores the fact that credence-forming processes should be sensitive to any germane evidence available to us.

One may think that there is something right about credences that are produced by a calibrated process even if such a process is insensitive to germane evidence that comes our way. For example, suppose we know that the long-run frequency of rain is 0.2. Based on such knowledge, we assign a credence of 0.2 to the

¹³See Murphy and Winkler 1977, pp. 41–2; DeGroot and Fienberg 1982, p. 14; and Seidenfeld 1985, p. 278.

proposition that it will rain on Wednesday. Such a credence may well be justified. But suppose that new evidence that points strongly in favour of the proposition's being true comes our way and we ignore it. Our credence of 0.2 now seems unjustified. One may, however, have the intuition that there is still something right about it. At least, assigning a credence of 0.2 to the proposition that it will rain on Wednesday still seems better than assigning a credence of 0.1 to it (for example).

Let us grant the intuition. In fact, we can account for it by holding that, all other things being equal, there is something epistemically good about having calibrated credence-forming processes. But calibration's being an epistemic good is compatible with its being insufficient for justification. Perhaps one might think that there are different senses of the term 'justification' and that there is a sense in which a credence is justified so long as it is produced by a calibrated process (especially in a context with lax epistemic standards).¹⁴ But even if so, there remains a good sense in which our credence of 0.2 is no longer justified—and (Calibration) cannot account for this. Furthermore, the sense in which our credence of 0.2 is not justified seems to be an important one that epistemologists care—or should care—about. After all, epistemologists typically hold that justification is defeasible. For instance, a justified belief that p may be defeated when we receive new evidence that tells against p . Similarly, the present case seems to be one in which our credence of 0.2 was once defeasibly justified but has now been defeated by new evidence. In contrast, any sense in which our credence of 0.2 remains justified even in the face of defeating evidence is not one in which justification is defeasible.

(Brier) avoids the above problem facing (Calibration). Assigning a constant credence of 0.2 to rain will yield a higher (that is, worse) average Brier score than assigning a higher credence to rain when it is going to rain and a lower credence to rain when it is not. Also, unlike (Probability), (Brier) does not rule out the possibility of an omniscient agent who justifiedly assigns credence 1 to all truths and credence 0 to all falsehoods. For according to it, a process produces justified

¹⁴Thanks to an anonymous referee for pressing me on this point.

credences so long as the credences it produces have a low average Brier score, whether or not they approximate the corresponding objective probabilities. In fact, according to (Brier), since the average Brier score of the credences formed by the omniscient agent equals 0, they are perfectly justified.

But (Brier) faces a problem of its own. Given (Brier), so long as a process produces an intermediate credence, *none* of the credences it produces will be perfectly justified. For once a process produces an intermediate credence, its average Brier score (or more accurately, that of the credences it produces) will be greater than 0. Given that justification comes in degrees and that the higher the average Brier score of the credences produced by a process, the less justified they are, any credence produced by a process with a positive average Brier score is less than perfectly justified.

However, intermediate credences may well be perfectly justified. According to Goldman, a belief-forming process is ‘a *functional operation* or procedure, i.e., something that generates a *mapping* from certain states—“inputs”—into other states—“outputs”’ (Goldman 1979, p. 11, his emphases). There is nothing wrong with an operation or procedure that leads to an intermediate credence if the input is unsuitable for generating a credence of 1 or 0. For example, suppose the input is a visual experience as of a silhouette of a horse and the output is an intermediate credence in there being a horse in front of us. Given the input, the forming of an intermediate credence seems to be in perfect order.

Lam holds that ‘Brier scoring captures the intuition that perfectly omniscient subjects, or subjects with only binary beliefs who have all true beliefs and no false beliefs, are a kind of epistemic ideal’ (Lam 2011, p. 217). But even if such subjects are an epistemic ideal—or even the only epistemic ideal—it does not follow that a subject’s doxastic states are less than perfectly justified if the ideal is not attained. The following analogy will help illustrate the point.

In decision theory, there is a distinction between the most *rational* choice and the *best* choice, where the most rational choice *maximizes expected utility*, whereas

the best choice *leads to the highest utility*. It is clear that the most rational choice is not always the best. By sheer misfortune, the most rational agent may not end up the most well off. But this is compatible with the ideal situation being one in which every choice we make in fact leads to the best outcome. Perhaps all we ultimately care about is maximizing our utility, but given that, the most rational thing to do is to act so as to maximize expected utility.

Similarly, we may grant that truth is the ultimate epistemic value and that it is an epistemic ideal that agents assign credence 1 to all truths and credence 0 to all falsehoods. This does not mean that a reliability theory of justified credence should deem any intermediate credence less than perfectly justified. In fact, analogously to the decision theory case, we should expect that intermediate credences may be perfectly justified even when the purported ideal is not met.¹⁵

At this point, a proponent of (Brier) may suggest that a credence produced by a process is perfectly justified if and only if the process has the lowest average Brier score among all *available* processes. This may allow processes with a positive average Brier score to produce intermediate credences that are perfectly justified—a process that assigns only credence 1 to truths and credence 0 to falsehoods may not be available.

But it is hard to say what counts as available in a way that will resolve the above worry for (Brier). It will not do to hold that a process counts as available just in case it is logically, physically, or biologically possible for an agent to form credences via the process. For a process that assigns only credence 1 to truths and credence 0 to falsehoods seems to be logically, as well as physically and biologically, possible. Also, given any process that has a positive average Brier score, nothing in the laws of logic, physics, or biology seems to rule out another process that has a slightly lower score. This makes it *too hard* for intermediate credences to be perfectly justified.

¹⁵Also, as Goldman points out, ‘[t]ruth is the main institutional goal of adjudication’, but ‘[s]ince evidential justification is the only reliable route to truth, [there is] an obligation on jurors to judge in accordance with the evidence’ (Goldman 2001, p. 40). Presumably, perfectly fulfilling such an obligation need not guarantee that one has arrived at the truth.

Now, there is a good sense in which an agent may lack the ability to do something that is logically, physically, or biologically possible for her to do. For example, I lack the ability to run 100 metres in less than 11 seconds, but such a feat is neither logically, physically, nor biologically impossible. Given this sense of ‘ability’, it is natural to suggest that a process is available to an agent if and only if she has the ability to form credences via the process. This may help make room for intermediate credences that are perfectly justified: an agent like God has the ability to assign nothing but credence 1 to truths and credence 0 to falsehoods, but most agents lack such an ability.

However, understanding ‘available’ in the way above makes it *too easy* for credences to be perfectly justified. Suppose an agent forms a credence via process m , which has a higher average Brier score than process n . Even if the agent’s cognitive limitations allow her to form credences via process m only, we will not want to say that her credence is thereby perfectly justified. For instance, consider an agent who knows that $\neg q$ and that q if p , but because of her inability to reason by *modus tollens*, assigns a relatively low credence to $\neg p$ via some other process available to her. Even if such a process has the lowest average Brier score among all the other processes available to the agent, her credence in $\neg p$ does not seem to be perfectly justified.

Perhaps one may suggest that a process is available to an agent if and only if some members of her community have the ability to form credences via the process. The agent in the example above may not have the ability to reason by *modus tollens*, but if some members of her community have the ability, the process of reasoning by *modus tollens* may still count as available. The problem with the suggestion, however, is that even if every member of the agent’s community lacks the ability in question, her credence in $\neg p$ still does not seem to be perfectly justified.

In sum, if a process counts as available just in case it is logically, physically, or biologically possible for an agent to form credences via the process, then it is too hard for intermediate credences to be perfectly justified. But if a process counts as

available if and only if an agent has the ability to form credences via the process, then it is too easy for intermediate credences to be perfectly justified. The worry that (Brier) makes no room for intermediate credences that are perfectly justified remains unresolved.¹⁶

4 Reliabilism and evidence

Process reliabilism, as traditionally formulated, cannot account for justified credence. And modified versions of the theory, such as (Probability), (Calibration), and (Brier), face certain worries. These worries, I will argue, stem from the same root: the theories falter because they pay scant attention to the evidence upon which our credences are based. Acknowledging this paves the way for my favoured version of process reliabilism—a hybrid theory inspired by William Alston that explicitly takes evidence into consideration (Alston 1988, 2005).

As Hume says, ‘[a] wise man . . . proportions his belief to the evidence’ (Hume 1748, p. 170). We may extract from this adage the plausible—and perhaps even platitudinous—idea that the strength of our credence ought to depend on the strength of the evidence on which it is based. But (Probability), (Calibration), and (Brier) pay hardly a nod to Hume’s adage. And that is why they falter. (Probability) falters when it penalizes us for assigning a credence of 0.6 to each proposition asserted by Jack, ignoring the fact that such an assignment is based appropriately on what we know about him. (Calibration) falters when it fails to penalize calibrated processes that assign credences based only on known long-run frequencies and that neglect any other germane evidence we possess. And (Brier) falters when it penalizes the assignment of intermediate credences, ignoring the fact that such an assignment may be perfectly justified when our evidence calls for it.

In contrast, the idea that our credences ought to reflect our evidence finds a natural home in *evidentialism*—one of process reliabilism’s main rivals. To account

¹⁶In fn. 34, after I have discussed a version of process reliabilism that incorporates evidentialist elements, I will look at another attempt to cash out ‘availability’ in terms of the evidence an agent possesses.

for justified belief, evidentialists do not appeal to reliability. Instead, they appeal to the idea that our doxastic attitudes should *fit* the evidence upon which they are based. According to Feldman and Conee, the main proponents of evidentialism,

S's doxastic attitude D at t toward proposition p is well-founded [or doxastically justified] 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 toward p fits e ; and
 - (c) there is no more inclusive body of evidence e' had by S at t such that having D toward p does not fit e' . (Feldman and Conee 1985, p. 24)

Now, we may question whether evidentialists can give us a non-circular and illuminating account of fit (Comesaña 2010, pp. 574–6). But let us grant that they can do so. Then, given that stronger evidence for p fits a higher credence in p , whereas weaker evidence for p fits a lower credence in p , evidentialism seems to embody Hume's adage.

Given that evidentialism seems well-suited to account for justified credence, whereas process reliabilism seems to have difficulty doing so, there may be some reason to steer towards the former. But the reliabilist need not yet despair. For all that has been said, a version of process reliabilism that incorporates evidentialist elements—in particular, one that gives us a reliabilist notion of fit—may still save the day. Comesaña (2010) has recently advanced a hybrid theory of reliabilism that individuates belief-forming processes according to the evidence on which beliefs are based. He does so in part to solve certain problems that traditional process reliabilism faces, such as the generality problem (of which more later) and the

problem posed by BonJour's clairvoyance examples.¹⁷ Goldman (2011) has also recently proposed a version of reliabilism that incorporates the notion of fit. He thinks that such an incorporation is a 'salutary addition' to process reliabilism and that its 'attraction . . . seems especially obvious in inferential justification', where one doxastic state is justified on the basis of another justified doxastic state (Goldman 2011, p. 263). In particular, Goldman 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.*).¹⁸ But prior to Comesaña (2010) and Goldman (2011), Alston has also advanced a version of reliabilism that gives centre stage to the evidence or *grounds* on which our beliefs are based (Alston 1988, 2005).¹⁹ And it is Alston's theory to which I now turn to help us account for justified credence.

4.1 Alston's theory of justified belief

According to Alston, a justified belief is one that is based on an adequate ground, where such a ground is adequate if and only if the objective probability of the belief's being true given that it is based on that ground is very high (Alston 2005, p. 99). A few things to note. First, for brevity, I will sometimes let talk about the probability of a belief's being true *given that it is based on a certain ground* be abbreviated by talk about the probability of a belief's being true *given that ground*. Second, on pain of circularity, a theory of doxastic justification that appeals to the notion of evidence or ground had better not cash out the notion as that which justifies one's belief or credence (Goldman 2011, p. 255). And neither Alston,

¹⁷The examples put forward by BonJour (1980, pp. 59–61) involve cases in which a person has reliable clairvoyant powers, uses such powers to form a belief that p , but has no evidence or reasons to believe that p or to think that she has clairvoyant powers. The worry is that the belief that p seems unjustified, but reliabilism says that it is justified.

¹⁸Goldman (2011, pp. 263–4) notes that mere fit is not sufficient for doxastic justification. To use his example, suppose Shirley makes a wild guess about the degree of confirmation of p upon q , guesses correctly, and thereby assigns the corresponding credence to q . Shirley's credence in q may fit her evidence p , but it does not seem justified. In light of this, Goldman requires that for Shirley's credence in q to be justified, it has to be formed by the right kind of process. But he does not go on to say much about how to evaluate Shirley's *credence*-forming processes.

¹⁹Comesaña (2010, p. 581) acknowledges that his own theory 'owes a great deal' to Alston (1988)'s.

Goldman, Comesaña, nor Feldman and Conee cash it out in such a way. Instead, they all take evidence or grounds to involve some kind of mental or psychological state or process.²⁰ Third, Alston talks about beliefs being based on *grounds* rather than on *evidence* because he thinks that ‘evidence’ is ‘a term used variously, at least by philosophers’, and ‘too squishy to mark precisely the most important distinctions between items that can function as grounds of beliefs’ (Alston 2005, pp. 82–3). But for my purposes, I will use ‘ground’ and ‘evidence’ interchangeably, and there will be no harm in doing so. Fourth, whereas the early Alston (1988) is interested in defending a theory of justification, the later Alston (2005) has sworn off such a business and is instead more interested in exploring various epistemic desiderata, one of which is that our beliefs be formed reliably.²¹ However, we may still take what he says about reliability to help us formulate a theory of doxastic justification. And for convenience, I will continue speaking of *Alston’s theory of justified belief* or *Alston’s theory of doxastic justification*.

Now, Alston’s theory is usually thought to be a kind of *indicator* reliabilism, according to which the justifiedness of a belief depends on whether the ground on which it is based is adequate or—as Alston sometimes puts it—on whether the ground is ‘sufficiently *indicative* of the truth of the belief’ (Alston 1988, p. 269, my emphasis). Formulated as such, the theory makes no explicit mention of processes. Furthermore, Alston did initially intend his theory to be contrasted with process reliabilism (*ibid.*, p. 281, p. 283). But he has recently argued that his theory is both a kind of indicator reliabilism *and* a kind of process reliabilism. What follows is a brief and simplified version of the argument.

According to Alston, for a belief *b* to be based on a ground *g* is just for *b* to be the output of a belief-forming process with *g* as its input, where the input is ‘something psychological—some psychological state or process’ such as a belief, a

²⁰See Alston 2005, p. 83; Comesaña 2010, p. 581; Goldman 2011, p. 257; and Feldman and Conee 2001, p. 2.

²¹To clarify, Alston (2005) is neither interested in defending a theory of doxastic justification nor a theory of propositional justification. See fn. 1 for more on the distinction between the two kinds of justification. See Alston 2005 for more on his epistemic desiderata approach.

memory, or a visual experience (Alston 2005, p. 83, p. 122, pp. 133–4).²² Hence, g is sufficiently indicative of the truth of b —that is, the objective probability of b being true given that it is based on g is high—if and only if the objective probability of b being true given that b is the output of a belief-forming process m with g as its input is high. Since Alston takes the relevant notion of objective probability to be hypothetical relative frequency, it is natural to hold that g is sufficiently indicative of the truth of b if and only if the frequency of true belief outputs relative to cases similar to that in which m takes g as its input and produces b as its output is high. But this is to say that m is reliable (with respect to the relevant cases). As Alston writes, ‘reliability of process and reliability of indicator turn out to coincide’ (Alston 2005, p. 137).

4.2 *Credences and Grounds*

Alston’s theory of justified belief is a kind of process reliabilism. And it is one that explicitly takes into consideration the grounds (or evidence) on which our beliefs are based. But as it stands, it is applicable to binary beliefs only. How may we modify or extend the theory so that it applies to credences?

4.2.1 Adapting Alston’s theory to deal with justified credences According to Alston, a belief is justified if and only if it is based on a certain ground, and the objective probability of the belief’s being true given that ground is very high. Now suppose that belief comes in degrees. It is natural to think that the stronger (or weaker) the belief, the higher (or lower) the objective probability in question needs to be for the belief to be justified. In particular, to account for justified credence, it is natural to put forward the following version of Alston’s theory (which I will later tweak):

(Grounds- α) For any $x \in [0, 1]$, any proposition p , and any subject a , a ’s credence of x in p is justified if and only if it is based on some ground g , where the

²²To be more careful, the ground ‘for a belief is not what we might call the total concrete 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, p. 268). See also Alston 2005, pp. 126–7.

objective probability of the credence having a true content given (that it is based on) g approximates or equals x .²³

(Grounds- α) offers us a reliabilist notion of fit—for a credence of x to fit the ground on which it is based is for the relevant objective conditional probability to approximate or equal x (where what counts as *approximate* may well be vague). Following Alston (2005), I will take objective probabilities to be hypothetical relative frequencies. But why invoke frequencies instead of objective chances, propensities, or even logical probabilities?

The problem with invoking objective chances or propensities is as follows. On some prominent views, the chance or propensity of an event happening in a deterministic universe—or of it having happened in the past—is either 1 or 0 (depending on whether the event will happen or happened) (Lewis 1980, Schaffer 2007). If such views are correct, then (Grounds- α) will have the following implausible consequence: no intermediate credences formed in a deterministic universe—or no intermediate credences about past events—will be justified. No doubt, some have argued that there are intermediate chances or propensities in a deterministic universe (Loewer 2001, Hoefer 2007). But it is best if the truth of (Grounds- α) is not held hostage to which view of chance or propensity is correct.

What about logical probabilities? Comesaña (2010, p. 597) is sympathetic to cashing out reliability in terms of logical probability. And Goldman (2011, p. 263) appeals explicitly to what he calls ‘degrees of confirmation’ in presenting his hybrid theory of reliabilism. But what exactly are logical probabilities? Notoriously, the project of spelling out what they are has met with various objections, and it is best if (Grounds- α) is not held hostage to whether the objections can be answered.²⁴ Now, Comesaña claims that many of the objections ‘are not relevant to the *existence* of the relation of logical probability’ (Comesaña 2010, p. 597, his emphasis). But

²³For simplicity, I focus on cases in which the grounds on which one’s credences are based are non-doxastic in nature. But what if our grounds consist of doxastic states, for example, credences? We may handle this by adding a recursive clause to the right-hand side of (Grounds- α) to the effect that, in such a case, such credences must themselves be justified.

²⁴See Hájek 2012 for a discussion of some of these objections. Alston (2005, pp. 104–8) offers more reasons for why logical probabilities are to be eschewed in cashing out the notion of reliability.

suppose we do not do more to spell out what logical probabilities are. Then when we hold that the logical probability of p given q is high, we might just be thinking that we *rationaly* ought to assign a high credence to p given q , or that we would be *justified* in assigning a high credence to p given q . If *that* is what logical probabilities amount to, then they do exist. But on pain of circularity, we should not invoke such probabilities in our theory of doxastic justification.²⁵

Given that the relevant notion of objective probability is relative frequency, why invoke *hypothetical* rather than *actual* frequencies? Well, doing so allows us to avoid worries having to do with the assignment of credences to one-off events. Consider a world in which there is just one coin and it is going to be tossed just once. A credence of 0.5 in the coin landing heads based on our perceiving that the coin is symmetrical may well be justified. But the actual frequency of the coin landing heads relative to our evidence is either 1 or 0, depending on whether the coin will in fact land heads or tails. Appealing to hypothetical cases in which the coin is tossed many times—or in which many similar coins are tossed—allows us to get around such a problem.

But there remains the question of how to deal with justified credences regarding one-off and *essentially non-repeatable* events. Consider the Big Bang or Darwin's birth in 1809. It is hard to make sense of either event happening more than once—even hypothetically—and it is not immediately clear what to count as similar events. Hence, one may worry that even if it makes sense to talk about the relative frequencies of such events, they will have a value of either 1 or 0. But intermediate credences in the universe beginning with the Big Bang or in Darwin being born in 1809 may well be justified.

To determine whether an intermediate credence of x in q is justified, where q is about a one-off and essentially non-repeatable event happening, I suggest that

²⁵As pointed out earlier, one may worry that evidentialists have not done enough to cash out what they mean by 'fit' in a non-circular and illuminating way (Comesaña 2010, pp. 574–6). In light of this, it is interesting to note that if the project of spelling out what logical probabilities are succeeds (without invoking the notion of rationality or justification), then such an objection loses its force. For evidentialists may then hold that a belief fits the evidence on which it is based if and only if the logical probability of the belief's being true given the evidence is high.

the relevant reference class to consider should not just contain cases in which one forms a credence of x in q .²⁶ (Otherwise, the objective probability of q given the grounds on which the credence is based will be either 1 or 0.) Instead, it should include *similar* cases in which one forms a credence of x in *other* propositions, where what counts as similar will depend on the process by which one's credence of x in q is formed. Consider, for example, a case in which one forms a credence about a one-off event on the basis of testimonial evidence. Suppose Ida hears her biology teacher assert the proposition that Darwin was born in 1809, and on such a basis, she forms a credence of 0.9 in it. To determine whether Ida's credence is justified, the relevant reference class should not just contain cases in which Ida assigns a credence of 0.9 to Darwin being born in 1809. Instead, it should contain cases in which she assigns a credence of 0.9 to other propositions, for example, cases in which she forms a credence of 0.9 in a proposition about Darwin on the basis of hearing her biology teacher assert it (whether such a proposition concerns Darwin's date of birth or other things about him).²⁷ If, in about 90% of such cases, Ida's credence of 0.9 is assigned to a true proposition, then her credence of 0.9 in Darwin being born in 1809 may well be justified.

At this point, one may wonder whether (Grounds- α) succumbs to the *generality problem*. Briefly, the problem has to do with there seeming to be no non-arbitrary way to classify the type to which a process token belongs in order to determine its reliability—classify it one way, and the process may be considered

²⁶I depart from Alston (2005) when he holds that if a process produces a belief with a certain content, then the type to which the process belongs must include only processes that produce beliefs with the same (and no other) content. As he puts it,

If the belief the formation of which we are considering is a belief that there is a maple tree in front of me, that ties down the output side, and there is no scope for choosing different ways of identifying it. (Alston 2005, p. 139)

Now, one motivation for such a position is that it helps us tame the generality problem (of which more later) by narrowing down the possible types to which a process belongs. But though I agree that the generality problem is to be reckoned with, such a way of doing so makes it hard for us to account for justified credences (and beliefs) regarding non-repeatable events.

²⁷Assume for simplicity that all of Ida's evidence for such propositions comes from the testimony of her teacher.

reliable; classify it another way, and it may not.²⁸ For instance, consider the process by which Ida, on hearing her teacher assert the proposition that Darwin was born in 1809, forms a credence of 0.9 in it. According to (Grounds- α), whether Ida's credence is justified depends on whether it matches the relevant relative frequency. This depends on what the relevant reference class is, which in turn depends on the type to which Ida's credence-forming process belongs. But one may worry that such a process belongs to indefinitely many types. Suppose that Ida's teacher is a human being, is a biology teacher, and has a low-pitched voice. Then Ida's credence-forming process appears to fall under each of the following types: *forming a credence of 0.9 in a proposition based on hearing a human being assert it*, *forming a credence of 0.9 in a proposition based on hearing a teacher assert it*, *forming a credence of 0.9 in a proposition about Darwin based on hearing a biology teacher assert it*, *forming a credence of 0.9 in a proposition about Darwin based on hearing Ida's biology teacher assert it*, *forming a credence of 0.9 in a proposition about Darwin based on hearing Ida's biology teacher assert it in a low-pitched voice*, etc.

If Ida's credence-forming process falls under indefinitely many types, which one is relevant for the purpose of evaluating whether her credence is justified? If the relevant type is *forming a credence of 0.9 in a proposition based on hearing a human being assert it*, then the corresponding reference class to consider will consist of all and only cases in which Ida forms a credence of 0.9 in a proposition based on hearing a human being assert it (regardless of whether the human being is a teacher and regardless of what proposition is asserted). If the relevant type is *forming a credence of 0.9 in a proposition about Darwin based on hearing a teacher assert it*, then the corresponding reference class will consist of all and only cases in which Ida forms a credence of 0.9 in a proposition about Darwin based on hearing a teacher assert it (regardless of whether the teacher is a biology teacher and regardless of whether the teacher is Ida's teacher). Since different reference classes will yield different relative frequencies, whether Ida's credence is justified or not depends

²⁸The problem was noted by Goldman (1979); for more on the problem, see Conee and Feldman 1998.

on which type is relevant. However, if there is no non-arbitrary way to pick out the relevant type, then it seems that there will be no non-arbitrary answer to the question of whether Ida's credence is justified.

How should a proponent of (Grounds- α) respond to the generality problem? A full solution to it is beyond the scope of this paper. After all, my aim is to come up with a way for reliabilists to account for justified credence; it is not to solve all the problems that traditional reliabilism faces. If, for instance, you have already decided that the generality problem is a fatal problem for reliability accounts of justified belief, it is unlikely that you will be sympathetic to such an account of justified credence. But if you are a reliabilist and think (or hope) that the problem can be solved, there is the additional issue of how we may extend reliabilism to deal with justified credence. And it is *this* issue that concerns me. Nevertheless, by drawing from the work of Alston (1995), let me at least allay the worry that (Grounds- α) will crumble in the face of the problem.

As Comesaña points out, many have tried to solve the problem by holding that 'the relevant process-type to assess for reliability must be chosen taking into account the *actual* production of the belief' (Comesaña 2006, p. 34, my emphasis). And Alston has offered a solution along such a line. Recall that, according to Goldman, a belief-forming process is 'a *functional operation* or procedure, i.e., something that generates a *mapping* from certain states—"inputs"—into other states—"outputs"' (Goldman 1979, p. 11; his emphases). Appealing to Goldman's understanding of a belief-forming process, Alston writes:

[E]very belief formation involves the activation of a certain psychologically realized *function*. That activation yields a belief with a propositional content that is a certain function (the function the psychological realization of which is activated here) of the proximate input. The function involved will determine both what features of the input have a bearing on the belief output and what bearing they have, i.e., how the content of the belief is determined by those features. (Alston 1995, p.

12, his emphasis)

The rough idea is that a psychological mechanism is activated in the formation of a doxastic attitude. Furthermore, this mechanism ‘embodies an input–output function’—it is the kind of mechanism that will take in a certain kind of input to produce a certain kind of output (*ibid.*, p. 17). And the kind of input the mechanism accepts plus the kind of output it produces in response to such input will determine the relevant process type. Or as Alston puts it, ‘[t]he function defines the epistemologically relevant type’ (*ibid.*).

Consider Ida’s credence in the proposition that Darwin was born in 1809. There are several possible mechanisms that could have been activated to produce such a credence. For example, perhaps the mechanism activated is one that takes as input all and only cases in which Ida hears a human being assert a proposition, and produces as its output a credence of 0.9 in it. Or perhaps the mechanism activated is one that takes as input all and only cases in which Ida hears her biology teacher assert a proposition about Darwin in a low-pitched voice, and produces as its output a credence of 0.9 in it. But suppose that the mechanism that is *in fact* activated is one that takes as input all and only cases in which Ida hears her biology teacher assert a proposition about Darwin, and produces as its output a credence of 0.9 in it. Then, to determine whether Ida’s credence is justified, the relevant reference class to consider will consist of all and only such cases.

To illustrate, suppose that Ida’s teacher happens to be a human being with a low-pitched voice. The relevant reference class will not consist of cases in which Ida forms a credence of 0.9 in a proposition based on hearing her father assert it (assuming that her father is not also her biology teacher). This is because the kind of psychological mechanism activated will accept as input *only* cases in which Ida hears her biology teacher—and not just any human being—assert a proposition about Darwin. Perhaps, if Ida had heard her father instead of her biology teacher assert that Darwin was born in 1809, she might have assigned a credence of 0.6 to the proposition. But in that case, some other psychological mechanism would have

been activated in the formation of such a credence. Also, the relevant reference class will not exclude cases in which Ida forms a credence of 0.9 in a proposition about Darwin based on hearing her biology teacher assert it in a high-pitched voice. This is because the kind of psychological mechanism activated will accept as input *all* cases in which Ida hears her biology teacher assert a proposition about Darwin—whether in a low-pitched voice or not.

There is no doubt more to be said about the above response to the generality problem.²⁹ For one thing, Alston admits that the solution he offers assumes the truth of ‘psychological realism’, that is, it assumes that, for any belief generated, there is a determinate psychological mechanism that produced it, where this mechanism is one that produces certain kinds of output in response to certain kinds of input (Alston 1995, pp. 18–19). This is an assumption he does not defend in full.³⁰ But as mentioned, a full solution to the generality problem is beyond the scope of this paper. Furthermore, there is a ‘partners in crime’ response available to us. First, it is well known that reliabilists in general have to deal with the problem (Conee and Feldman 1998). Second, the other accounts of justified credence discussed earlier also face the task of telling us how to pick out the relevant process type. This is because they are all versions of *process* reliabilism. And whether a process tends to produce credences that match the corresponding objective probabilities, whether it is calibrated, or whether it produces credences with a low average Brier score depends on the type to which it belongs. Of course, the preceding does not show that the generality problem is not a problem. But because it is a problem for reliabilists in general, the lack of a full solution to it poses no particular embarrassment to (Grounds- α).³¹

²⁹For more discussion on Alston’s solution, see Adler and Levin 2002, Feldman and Conee 2002, and Comesaña 2006.

³⁰He does, however, offer some considerations in favour of psychological realism in Alston 1995, pp. 19–21.

³¹Further, it is worth noting that Comesaña (2006) has argued that any complete epistemological theory of justification, including evidentialism, has to deal with some version of the generality problem.

4.2.2 A problem with (Grounds- α) and a tweak We have seen how the proponent of (Grounds- α) may respond to the generality problem. But as it stands, the theory still needs to be tweaked. Suppose we assign a high credence to there being an elephant in front of us based on our visual experience as of there being an elephant in front of us. And suppose that the objective probability of there being an elephant in front of us given that our credence is based on such an experience is high. Our credence may none the less be unjustified if, say, we have ignored our memory of having taken a drug that causes hallucinations of elephants.

In response, one may hold that a credence of x in p is justified only if it is based on all relevant and available grounds.³² But what counts as relevant? Intuitively, a ground is relevant to a belief if it bears on the justifiedness of the belief. But spelling out ‘relevant’ in such a way leads to circularity. And holding that a credence is justified only if based on all available grounds—whether relevant or not—is unduly demanding. It is not reasonable, for instance, to require that our credences be based on *all* of our memorial beliefs.

Fortunately, we may avoid talk of relevance and solve the problem by taking a leaf from Feldman and Conee’s account of well-foundedness or doxastic justification, which we encountered at the beginning of section 4. Clause (c) of condition (ii) in their statement of well-foundedness is the most germane to us. Taking a leaf from them, we may amend (Grounds- α) to get:

(Grounds) For any $x \in [0, 1]$, any proposition p , and any subject a , a ’s credence of x in p is justified if and only if

1. it is based on some ground g , where the objective probability of the credence having a true content given (that it is based on) g approximates or equals x , and
2. there is no more inclusive ground g' had by a such that the objective probability of the credence having a true content given (that it is based

³²Cf. Alston 2005, pp. 100–3.

on) g' neither approximates nor equals x .

(Grounds) avoids the problem above. Suppose we form a high credence in there being an elephant in front of us, ignoring our memory of having taken a drug that induces hallucinations of elephants. According to (Grounds), our high credence is unjustified. For we have a more inclusive ground g' such that the objective probability of there being an elephant in front of us given g' is low.

4.2.3 (Grounds) versus other reliability accounts of justified credence (Grounds)

also avoids the problems faced by the other reliability accounts of justified credence discussed earlier. First, let us compare (Grounds) with the account discussed at the end of section 2, according to which a credence is justified if and only if it is produced by a process that tends to lead to a high proportion of credences with true contents. As we have seen, the latter account is too indiscriminating. Suppose we know that the frequency of rain in our city is 90%, but we assign a credence of 0.51 to rain on a particular day. The account cannot explain why our credence is unjustified, or at least, less justified than a credence of 0.9 in rain. But (Grounds) can do so. Furthermore, (Grounds) can also explain why a credence of 0.51 in rain, though unjustified *tout court*, still seems more justified than a credence of 0.2 in rain. For given that justification comes in degrees, we may hold that the closer our credences approach the corresponding objective probabilities, the more justified they are.³³

Second, let us compare (Grounds) with (Probability). Unlike (Probability), (Grounds) does not rule out the possibility of an omniscient agent who justifiedly assigns credence 1 to all truths (and credence 0 to all falsehoods). For the agent's credences are presumably based on perfectly adequate grounds such that for any proposition p and any ground g , the objective probability of p given g equals 1. In such a case, even if the unconditional objective probability of p is of intermediate

³³How should we measure degrees of justification? Suppose our credence in rain is 0.8 and the relevant objective probability is 0.9. To what degree is our credence justified? Also, how much more justified is a credence of 0.8 in rain than a credence of 0.7 in rain (for example)? These are interesting questions, but I am afraid that I cannot do full justice to them here.

value, a credence of 1 in p is justified (given there is no more inclusive ground such that the objective probability of p given such a ground is less than 1).

Recall also the example in which Jack asserts a number of propositions, each of which we initially had no evidence whatsoever to think is true or false. If we know that Jack asserts something true 60% of the time and something false 40% of the time, a credence of 0.6 in each proposition asserted seems justified. And this is so even if the objective probability of each proposition being true is not close to 0.6 at all. As we have seen, (Probability) penalizes such an assignment of credences. But (Grounds) does not. Consider the proposition that Jill is ill, and suppose that we had no evidence whatsoever as to its truth until Jack asserted it. According to (Grounds), a credence of 0.6 in the proposition is justified if based on the ground that Jack has just asserted it. For suppose we wish to compute the relevant relative frequency. The appropriate reference class plausibly consists of cases in which we assign a credence of 0.6 to a proposition (that we initially had no evidence to think is true or false) based on the ground that Jack has asserted it. And the frequency of us assigning a credence of 0.6 to a true proposition relative to such cases will presumably approximate 0.6 (given that Jack asserts true propositions 60% of the time).

Third, let us compare (Grounds) with (Calibration). Recall the weather forecaster who knows that the long-run frequency of rain in a particular city is 0.2. Suppose that, on this basis, she assigns a credence of 0.2 to rain each day while ignoring evidence having to do with the latest and most relevant meteorological data. (Calibration) does not penalize such an assignment of credences, because the relevant credence-forming process will be calibrated (at least with respect to the credences in question). Thus (Calibration) is too indiscriminating. But (Grounds) avoids this problem. For according to it, the weather forecaster's credences are unjustified if she has some more inclusive ground such that the relative frequency of rain given such a ground does not equal or approximate 0.2.

Finally, let us compare (Grounds) with (Brier). As we have seen, (Brier)

faces the worry that intermediate credences are never perfectly justified. But (Grounds) avoids such a worry—an intermediate credence of x based on g may be perfectly justified if the objective probability of the credence having a true content given that it is based on g equals x exactly.³⁴

In sum, by focusing on the grounds or evidence upon which our credences are based—by giving us a reliabilist notion of fit—(Grounds) allows us to account for justified credence while avoiding the problems faced by (Probability), (Calibration), and (Brier).

5 Conclusion

Traditional process reliabilism cannot account for justified credence. Do modified versions of the theory fare better? I have discussed three such theories—(Probability), (Calibration), and (Brier)—and argued that they face certain problems. I have also suggested that, because the theories pay scant attention to the evidence on which our credences are based, they do not account for justified credence well. In their stead, I have defended (Grounds), a version of Alston’s theory of doxastic justification. I have argued that, by weaving evidentialist elements into reliabilism, it accounts for justified credence better than the other theories discussed.

One might wonder if it is a concession to the evidentialist—or to internalists in general—for the reliabilist to subscribe to (Grounds). After all, Alston (1988) calls his theory *internalist externalism*, and the notion of grounds, as used by him, is

³⁴As we have seen, the proponent of (Brier) may suggest that an intermediate credence is perfectly justified if and only if the process that produced it has the lowest average Brier score among all the available processes. And as we have seen, it is hard to cash out ‘available’ in a way that neither makes it too hard nor too easy for intermediate credences to be perfectly justified. But perhaps the proponent of (Brier) may try to cash out ‘availability’ in terms of the evidence possessed by an agent. For instance, one may hold that a process that assigns credence 1 to all truths and credence 0 to all falsehoods is not available to ordinary agents, because the evidence such agents possess does not permit—or *fit*—such an assignment of credences.

Now, such a move is consistent with my larger point that a reliability theory of justified credence should not neglect the evidence upon which our credences are based. So it is all to the good if a version of (Brier) that explicitly takes an agent’s evidence into consideration can account for justified credence. It is hard, however, to see how (Brier) can on its own give us a reliabilist theory of fit. We may, of course, be able to appeal to some other means to provide such a theory. But then, instead of appealing to Brier scoring to account for justified credence, it will be much more straightforward to appeal to such means directly. And I have done precisely that by defending (Grounds).

an internalist one. However, if there is a concession, it is a harmless one. For according to (Grounds), whether a credence is justified depends on whether it matches the relevant objective probability. And this crucial condition for justification is clearly externalist and distinguishes (Grounds) from an internalist theory such as evidentialism.³⁵ Admittedly, (Grounds) seems to be an amendment of reliabilism that is friendlier to Alston than to Goldman, especially since the latter has in the past been reluctant to make any concession to internalism (e.g. in Goldman 1999a). But as mentioned, even Goldman has recently defended a hybrid reliabilism that incorporates evidentialist elements.³⁶ In so far as the reliabilist wishes to account for justified credence, that is the way to go.³⁷

³⁵By an *internalist* theory of justification, I mean one according to which a person's doxastic attitudes are 'justified only by things that are internal to the person's mental life' (Feldman and Conee 2001, p. 2). Accordingly, an externalist theory of justification is one according to which a person's doxastic attitudes are justified by things external to her mental life. Following Bergmann (1997), we may also distinguish between a strong version of externalism according to which all conditions for justification are external, and a more moderate version according to which, while some very important conditions for justification are external, there are some that are internal. (Grounds) is a moderate externalist theory of justification. (But note that Bergmann (1997) focuses on warrant rather than justification; also, he focuses on *access internalism* and *access externalism*, which have to do with whether a person has some sort of special access to that which justifies her doxastic attitudes.)

³⁶But Goldman (2011) sees his hybrid reliabilism as lending a helping hand to evidentialism. He thinks that 'evidentialism . . . could benefit by incorporating reliabilist themes' (Goldman 2011, p. 254).

³⁷This paper has benefited immensely from valuable feedback from Jens Christian Bjerring, Ben Blumson, Mark D'Cruz, Yongming Han, Michael Pelczar, Neil Sinhababu, and the audience at the 2012 Australasian Association of Philosophy Conference. Thanks also to two anonymous referees and the editor of *Mind* for very helpful comments.

References

- Adler, Jonathan and Michael Levin 2002: 'Is the Generality Problem too General?' *Philosophical and Phenomenological Research*, 65, pp. 87–97.
- Alston, William P. 1988: 'An Internalist Externalism'. *Synthese*, 74, pp. 265–283.
- 1995: 'How to Think about Reliability'. *Philosophical Topics*, 23, pp. 1–29.
- 2005: *Beyond "Justification": Dimensions of Epistemic Evaluation*. US: Cornell University Press.
- Bergmann, Michael 1997: 'Internalism, Externalism and the No-Defeater Condition'. *Synthese*, 110, pp. 399–417.
- BonJour, Laurence 1980: 'Externalist Theories of Empirical Knowledge'. *Midwest Studies in Philosophy*, 5, pp. 53–73.
- Brier, Glenn W. 1950: 'Verification of Forecasts Expressed in Terms of Probability'. *Monthly Weather Review*, 78, pp. 1–3.
- Cherniak, Christopher 1981: 'Minimal Rationality'. *Mind*, 90, pp. 161–183.
- Cohen, Robert S. and Larry Laudan (eds) 1983: *Physics, Philosophy and Psychoanalysis: Essays in Honor of Adolf Grünbaum*. Dordrecht: D. Reidel Publishing Co.
- Comesaña, Juan 2006: 'A Well-Founded Solution to the Generality Problem'. *Philosophical Studies*, 129, pp. 27–47.
- 2010: 'Evidentialist Reliabilism'. *Noûs*, 44, pp. 571–600.
- Conee, Earl and Richard Feldman 1998: 'The Generality Problem for Reliabilism'. *Philosophical Studies*, 89, pp. 1–29.
- DeGroot, Morris H. and Stephen E. Fienberg 1982: 'The Comparison and Evaluation of Forecasters'. *The Statistician*, 32, pp. 12–22.

- Dougherty, Trent (ed.) 2011: *Evidentialism and its Discontents*. New York: Oxford University Press.
- Fairweather, Abrol and Linda Zagzebski (eds) 2001: *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. New York: Oxford University Press.
- Feldman, Richard and Earl Conee 1985: 'Evidentialism'. *Philosophical Studies*, 48, pp. 15–34.
- 2001: 'Internalism Defended'. *American Philosophical Quarterly*, 38, pp. 1–18.
- 2002: 'Typing Problems'. *Philosophical and Phenomenological Research*, 65, pp. 98–105.
- Firth, Roderick 1978: 'Are Epistemic Concepts Reducible to Ethical Concepts?' In Goldman and Kim 1978, pp. 215–229.
- Goldman, Alvin I. 1979: 'What is Justified Belief?' In Pappas 1979, pp. 1–23.
- 1986: *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- 1988: 'Strong and Weak Justification'. *Philosophical Perspectives*, 2, pp. 51–69.
- 1999a: 'Internalism Exposed'. *Journal of Philosophy*, 96, pp. 271–293.
- 1999b: *Knowledge in a Social World*. New York: Oxford University Press.
- 2001: 'The Unity of the Epistemic Virtues'. In Fairweather and Zagzebski 2001, pp. 30–48.
- 2011: 'Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism's Problems, Reliabilism's Rescue Package'. In Dougherty 2011, pp. 254–280.
- 2012: *Reliabilism and Contemporary Epistemology*. New York: Oxford University Press.

- Goldman, Alvin I. and Jaegwon Kim (eds) 1978: *Values and Morals*. Dordrecht: D. Reidel Publishing Co.
- Hájek, Alan 2012: ‘Interpretations of Probability’. *The Stanford Encyclopedia of Philosophy* [online encyclopedia], Winter 2012 Edition, <<http://plato.stanford.edu/archives/win2012/entries/probability-interpre/>>, accessed 6 March 2013.
- MS: ‘A Puzzle About Partial Belief’.
- Hoefer, Carl 2007: ‘The Third Way on Objective Probability: A Sceptic’s Guide to Objective Chance’. *Mind*, 116, pp. 549–596.
- Hume, David 1748: *An Enquiry Concerning Human Understanding*, ed. Tom L. Beauchamp. New York: Oxford University Press, rpnt 1999.
- Jeffrey, Richard C. (ed.) 1980: *Studies in Inductive Logic and Probability: Volume II*. Berkeley and Los Angeles, CA: University of California Press.
- Joyce, James M. 1998: ‘A Nonpragmatic Vindication of Probabilism’. *Philosophy of Science*, 65, pp. 575–603.
- Kahneman, Daniel, Paul Slovic, and Amos Tversky (eds) 1982: *Judgment Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.
- Lam, Barry 2011: ‘On the Rationality of Belief-Invariance in Light of Peer Disagreement’. *The Philosophical Review*, 120, pp. 207–245.
- Lewis, David 1973: *Counterfactuals*. UK: Oxford, Basil Blackwell.
- 1980: ‘A Subjectivist’s Guide to Objective Chance’. In Jeffrey 1980, pp. 263–293.
- Lichtenstein, Sarah, Baruch Fischhoff, and Lawrence D. Phillips 1982: ‘Calibration of Probabilities: The State of the Art to 1980’. In Kahneman, Slovic, and Tversky 1982, pp. 306–334.

- Loewer, Barry 2001: 'Determinism and Chance'. *Studies in History and Philosophy of Modern Physics*, 32, pp. 609–620.
- Murphy, Allan H. and Robert L. Winkler 1977: 'Reliability of Subjective Probability Forecasts of Precipitation and Temperature'. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 26, pp. 41–47.
- Pappas, George S. (ed.) 1979: *Justification and Knowledge: New Studies in Epistemology*. Dordrecht: D. Reidel Publishing Co.
- Rissler, James 2006: 'Open Theism'. *Internet Encyclopedia of Philosophy* [online encyclopedia], <<http://www.iep.utm.edu/o-theism/>>, accessed 1 Feb 2014.
- Schaffer, Jonathan 2007: 'Deterministic Chance?' *British Journal for the Philosophy of Science*, 58, pp. 113–140.
- Seidenfeld, Teddy 1985: 'Calibration, Coherence, and Scoring Rules'. *Philosophy of Science*, 52, pp. 274–294.
- van Fraassen, Bas C. 1983: 'Calibration: A Frequency Justification for Personal Probability'. In Cohen and Laudan 1983, pp. 295–319.
- 1984: 'Belief and the Will'. *The Journal of Philosophy*, 81, pp. 235–256.