

Whither Evidentialist Reliabilism?

Juan Comesaña
University of Arizona

1. Introduction

Evidentialism and Reliabilism are two of the main contemporary theories of epistemic justification. Some authors have thought that the theories are not incompatible with each other, and that a hybrid theory which incorporates elements of both should be taken into account.¹ More recently, other authors have argued that the resulting theory is well-placed to deal with fine-grained doxastic attitudes (credences).²

In this paper I review the reasons for adopting this kind of hybrid theory, paying attention to the case of credences and the notion of probability involved in their treatment. I argue that the notion of probability in question can only be an epistemic (or evidential) kind of probability. I conclude that the resulting theory will be incompatible with Reliabilism in one important respect: it cannot deliver on the reductivist promise of Reliabilism. I also argue that attention to the justification of basic beliefs reveals limitations in the Evidentialist framework as well. The theory that results from the right combination of Evidentialism and Reliabilism, therefore, is neither Evidentialist nor Reliabilist.

2. Evidentialism

Evidentialism has been defined by Conee and Feldman (1985) as follows:

Evidentialism: Doxastic attitude *D* toward proposition *p* is epistemically justified for *S* at *t* if and only if having *D* toward *p* fits the evidence *S* has at *t*.

Three questions need to be answered before we have a full understanding of Evidentialism: what kinds of things can *be* evidence?; what is it for subject to *have* some evidence?; and what is it for a body of evidence to *fit* a doxastic attitude?

Conee and Feldman themselves have a conception of evidence and its possession according to which two subjects in the same total (non-factive) mental states cannot differ in what evidence they possess. But one can combine Evidentialism with other conceptions

¹See, for instance, Comesaña (2010a), Comesaña (2010b) and Goldman (2011). An important precursor is Alston (1988) (although Alston doesn't explicitly discuss Evidentialism).

²See Dunn (2015), Tang (2016b) and Pettigrew (ms).

of evidence and its possession and end up with a package of views which denies the supervenience of epistemic justification on (non-factive) mental states. For instance, if one thinks of evidence as consisting of true propositions or facts and its possession as consisting in knowledge, then two subjects can be in the same (non-factive) mental states and yet differ on what evidence they have.^{3,4} A third option is to say that evidence consists of propositions, and that experiences (even when non-veridical) provide these propositions as evidence.⁵ This is in fact the position I favor, and I will come back to it below.

One kind of mentalist Evidentialism (at least inspired by Conee and Feldman), then, would answer our three questions as follows: evidence is constituted by mental states, in particular by justified beliefs and experiences; a subject has some evidence just in case he is in the relevant mental state; and the fitting relation between bodies of evidence and doxastic attitudes is a primitive, non-reducible epistemic fact.⁶ Notice the apparent circularity in the answer to the first question: Evidentialism has it that justification supervenes on evidence, and we are told that justified beliefs can be evidence. This circularity is benign provided that one thinks justification has a recursive structure. In very rough terms, the idea is that experience (together perhaps with ostensible memories) provides us with *prima facie* non-inferential justification, and beliefs thus justified by experience can combine to produce further justified doxastic attitudes. I come back to this issue below.

Evidentialism is a theory of *propositional* justification—of what it is for a doxastic attitude to be justified for a subject, independently of whether the subject adopts that attitude. We also need a theory of *doxastic* justification—of what it takes for an attitude to be justifiedly adopted. It won't do just to say that an attitude is doxastically justified just in case it is propositionally justified and adopted: subjects may adopt the right attitudes for the wrong reasons. Conee and Feldman themselves propose the following theory of “well-foundedness” to add to their Evidentialism:

³The equation of evidence with knowledge is from Williamson (2000), and the parenthetical regarding non-factive mental states is designed to make room for Williamson's own conception of knowledge as itself a mental state. In what follows I omit the qualification.

⁴Supervenience may be too weak a notion to capture the essence of Evidentialism. The traditional definition of supervenience, applied to our case, is simply that there cannot be a difference as to what attitudes are justified for some subjects without there being a difference as to which evidence those subjects have. Combined with a mentalist conception of evidence and its possession, this yields the further supervenience thesis to the effect that there cannot be a difference as to which attitudes are justified for some subjects without a difference in the mental states they are in. But, plausibly, the Evidentialist and the Mentalist want more than mere supervenience: they may want not just the existence of a mere co-variation, but a *constitutive* relation between justification and evidence. If it turns out, say, that justification and mental states co-vary in the requisite way only because they in turn co-vary with a third condition, the resulting view need not be particularly friendly to Evidentialism. An analogy may help bring the point home. Suppose that we define Physicalism as the thesis that every fact supervenes on physical facts. That thesis is compatible with Cartesian substance dualism, as long as the non-physical stuff exists necessarily. Maybe the supervenience thesis is interesting in its own right, but conceiving of physicalism as compatible with substance dualism does not get the spirit of the view right. Analogously, one would have thought that Evidentialism would have to be incompatible with non-evidential facts determining epistemic justification, even when they obtain necessarily.

⁵For a development of this view, see Comesaña and McGrath (2014) and Comesaña and McGrath (2016).

⁶Conee and Feldman themselves hold that there is no such thing as unpossessed evidence. They seem to think of evidence, then, as *token*, instantiated mental states, rather than uninstantiated types—see Conee and Feldman (2008).

S's doxastic attitude D at t toward proposition p is well-founded if and only if:

- (i) having D toward p is justified for S at t; and
- (ii) S has D toward p on the basis of some body of evidence e, such that
 - (a) S has e as evidence at t;
 - (b) having D 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'.

Notice that Conee and Feldman are relying here on the notion of basing an attitude on a body of evidence. The Evidentialist notion of well-foundedness is thus importantly different from the Evidentialist notion of justification. To see the difference, consider the Williamsonian theory of evidence briefly alluded to earlier: items of evidence are facts, and they are possessed by a subject when they are known by that subject. That Williamsonian theory of evidence, as we said before, is compatible with at least the letter of Evidentialism. On the Williamsonian view, a proposition can be justified for a subject by either being sufficiently supported by the evidence the subject has or by being part of that evidence (in which case it trivially counts as being sufficiently supported by the evidence).⁷ There will therefore be no difference in the justificatory status of any proposition for any subjects without a difference in the evidence possessed by those subjects. Suppose, however, that we add to that kind of the theory the claim that one can be justified in believing the propositions which are part of one's evidence even if one does not base that belief on any evidence. Indeed, barring controversial cases of higher-order evidence, justification for believing the propositions which are part of one's evidence will in general require that one *not* base those beliefs on any evidence. Rather, those basic beliefs will be the result of direct knowledge by different modalities, for instance by looking. When you know that there is a snowball in front of you because you see it, it is part of your evidence that there is a snowball in front of you, and you are justified in believing that there is a snowball in front of you, but your belief that there is a snowball in front of you is not based on any evidence you have (your belief is certainly not based on itself, and you may have no other relevant evidence). The resulting view, however, is not compatible with the well-foundedness theory of basing, for that theory requires that all of one's justified beliefs be based on evidence. I return to this important issue below.

Why the need for clause (ii)(c)? Conee and Feldman's idea here is that even if the part of the subject's evidence on which he bases his attitude does indeed fit that attitude, there may be other parts of his evidence which don't. Suppose, for instance, that I believe that Fred is older than 9, and that I base this belief on the evidence that I have that Fred has gray hair. Compatible with all that, it may also be part of my evidence that Fred suffers from a condition that may cause premature greying of the hair. Moreover, what matters is which attitude my *total* evidence justifies—for it may also be part of my evidence that

⁷Williamson has recently added to his epistemology the claim that a body of evidence fully justifies a proposition only if it entails it—see Williamson (2013) and Williamson (forthcoming), and cf. Cohen and Comesaña (2013), Cohen and Comesaña (forthcoming) and Comesaña (2017).

Fred goes to college. In that case, my belief would be well-founded only if based on those three relevant pieces of evidence, and it wouldn't be if based only on the color of Fred's hair.⁸

One advantage of evidentialism and the accompanying notion of well-foundedness is that it applies to doxastic attitudes in general, and not just to beliefs. Thus, our evidence can fit disbeliefs and suspensions of judgments as well as beliefs, and they may also fit degrees of beliefs (credences).

One can have misgivings about different aspects of this kind of Evidentialism. The main worry that I am interested in now focuses on its primitivism, on the fact that it doesn't have much to say about why certain bodies of evidence "fit" certain doxastic attitudes. Reliabilism promises to deliver on precisely that front, but it faces problems of its own.

3. Reliabilism

Goldman (1979) is responsible for establishing Reliabilism as a theory of epistemic justification. A rough version of such Reliabilism has it that a belief is justified if and only if it is produced by a reliable belief-forming process. A bit less roughly, a belief is justified if and only if it is produced by a belief-independent belief-forming process and that process is reliable, or it is produced by a belief-dependent belief-forming process (i.e., a process some of whose inputs are the contents of some of the subject's beliefs) and that process is conditionally reliable (i.e., it tends to produce true beliefs given that its belief-dependent inputs are true).

Such Reliabilism is subject to three different kinds of objections. First, the objection that reliability is not necessary for justification. Second, the objection that reliability is not sufficient for justification. Third, the objection that we have no principled way of measuring the reliability of a belief-forming process. The first objection arises from Cohen's "new evil demon" problem (Cohen (1984)), the second from Bonjour's clairvoyant cases (Bonjour (1980)), and the third from Conee and Feldman's "generality problem" (Conee and Feldman (1998)—the problem was already noticed by Goldman (1979)).

The generality problem is more fundamental than the other two, because even formulating the other two problems presupposes an answer to the generality problem. The generality problem arises from the fact that any token belief-forming process will belong to an indefinite number of belief-forming process-types. It is usually assumed that reliability can only be defined for types, because types (but not tokens) are repeatable.⁹ But if any token belongs to indefinitely many types, and if those types differ in reliability, then we need a principled way to select a type for each token. For instance, if I believe that it is precipitating based on my belief that it is raining, then the token process instantiates

⁸Depending on one's account of the basing relation, one may hold that the belief need not be based on all three items of evidence. In any case, the main point is that the belief will not be justified if based only on the fact that Fred has gray hairs.

⁹Comesaña (2006) argues against this assumption that we can make sense of the reliability of a token process, but also notes that this will not help Reliabilists avoid the generality problem.

the type of believing a proposition on the basis of another proposition which entails it, but also the type of believing a proposition about the weather on the basis of another proposition about the weather, and the reliability of those two types differ substantially.

Once we have a solution to the generality problem (that is to say, once we have a principled way of associating each token belief-forming process with a type to be assessed for reliability) we need to figure out how the reliability of that type is evaluated. This will be the task of a later section.

Cohen's new evil-demon problem and Bonjour's clairvoyant problem assume that we have a solution to the generality problem and an answer to the question of how to assess the reliability of belief-forming process-types. In particular, Cohen's new evil-demon problem assumes that the beliefs of a victim of an evil-demon are not produced by reliable belief-forming processes, and, given that they are justified, concludes that reliability is not necessary for justification. Bonjour's clairvoyant problem assumes that the beliefs of a genuine clairvoyant are reliably produced, and, given that they are not justified, concludes that reliability is not sufficient for justification.

A different kind of objection is that whereas Reliabilism as developed can account (perhaps) for the justification of beliefs, it is not clear how to adapt it to other doxastic attitudes. Within the realm of coarse-grained epistemology, we could try the following: just as a belief is justified if the process that produced it is reliable, so too disbelief is justified if the process is anti-reliable (or, equivalently, if the process which produces the belief in the negation of the relevant proposition is reliable), and suspension of judgment is justified if the process is neither reliable nor anti-reliable. And within fine-grained epistemology, we could perhaps try the following: a degree of belief (or credence) is justified if and only if it matches the degree of reliability of the process that produced it. Those options might well work,¹⁰ but they raise the same issue we touched upon above: how to measure the reliability of a belief-forming process-type. Before tackling that issue, however, we need a solution to the generality problem.

4. How to Solve the Generality Problem

Alston (1988) proposes a version of reliabilism according to which a belief is justified just in case it is based on an adequate ground, where the adequacy of a ground is a matter of its reliability. In place of Alston's grounds, we can invoke the notion of evidence. This will give us a principled way of selecting a type for each belief-forming process-type: the type *believing that p on the basis of e*. Remember that Evidentialists themselves appeal to this notion of basing in their account of well-foundedness. It is *that type* which should be assessed for reliability. This is what I proposed in Comesaña (2006). My proposal was the following:

Well-Founded Reliabilism (first pass): A belief that *p* by *S* is epistemically justified if and only if:

¹⁰Although see Tang (2016a) for more on how Reliabilists should capture suspension of judgment.

- (i) *S* has evidence *E*;
- (ii) the belief that *p* by *S* is based on *E*; and
- (iii) the type producing a belief that *p* based on evidence *E* is a reliable type.

That proposal, however, ignores the reasoning behind clause (ii)(c) of the evidentialist definition of well-foundedness. Recall the reason: if I base my belief on a subset of my evidence which justifies it, but ignoring some other evidence which I have which does not justify it, then my belief is not well-founded. Taking that into account yields the following refined version of well-founded reliabilism:

Well-Founded Reliabilism: A belief that *p* by *S* is epistemically justified if and only if:

1. *S* has evidence *E*;
2. the belief that *p* by *S* is based on *E*;
3. the type *producing a belief that p based on evidence E* is a reliable type; and
4. there is no more inclusive body of evidence *E'* had by *S* at *t* such that the type *producing a belief that p based on evidence E'* is not a reliable type.

Notice that if we bracket clause 2, what is left is a definition of propositional justification for Reliabilism.

When assessing Evidentialism, we said that the theory needed an account of evidence and its possession. Well-founded Reliabilism inherits those needs, and the same options canvassed earlier are available here as well. In fact, however, a further development of Well-Founded Reliabilism will provide us with an argument for a specific account of evidence. I turn to that development next.

5. How to Measure Reliability

We have a solution to the generality problem insofar as we have a principled way of selecting a process-type for any token process of belief formation. But this doesn't yet give us a complete reliabilist theory, for we need to figure out how to measure the reliability of a process type.

One possible answer is to think of the reliability of a process as the truth-to-falsity ratio of its outputs. There are two varieties of this truth-ratio conception of reliability: reliability as actual high truth-ratio, or reliability as counterfactual high truth-ratio. The first one counts a process as reliable if and only if the ratio of truths to falsehoods within its actual outputs is sufficiently high, whereas the second one concerns not just actual truth-ratios but counterfactual ones as well. A counterfactual conception is preferable insofar

as there may be processes which, although intuitively justification-conferring, do not have many actual outputs, and so its actual truth-ratio may be coincidentally low. Going to a counterfactual account of reliability is not without issues, however, because one needs to restrict the relevant range of counterfactual applications in order to not trivialize the account. This issue with the counterfactual account is reminiscent of the generality problem, for we need to decide which of the counterfactual applications of the process are the relevant ones to measure reliability.

A different measure of reliability has also been proposed by Alston (and used in Comesaña (2009)): reliability as high conditional probability. The idea here is that a type of the form *believing that p based on e* is reliable if and only if, for some suitable probability function Pr and some suitable threshold r , $Pr(p|e) \geq r$. Two immediate questions about that approach are: what is a suitable probability function, and what is a suitable threshold? A third issue has to do with Carnap's distinction between confirmation as firmness and confirmation as increase in firmness.¹¹ $Pr(p|e)$ can be greater than r even if e doesn't raise the probability of p —indeed, for $r < 1$ it can happen that $r < Pr(p|e) < Pr(p)$. If this is the case, then we can hardly say that it is e which justifies the subject's belief that p . This need not be a problem for the Reliabilist who appeals to the notion of evidence only to solve the generality problem, and who doesn't take himself to be elucidating the notion of evidential justification, but only justification simpliciter. It may, however, be an issue for the Evidentialist—we'll come back to this point later.

Let us leave those questions aside for the moment, and formulate the resulting version of reliabilism:

Probabilistic Evidentialist Reliabilism: A belief that p by S is justified if and only if:

1. S has evidence E ;
2. the belief that p by S is based on E ;
3. $Pr(p|E) \geq r$.
4. There is no more inclusive body of evidence E' had by S such that $Pr(p|E') < r$.

The resulting theory is called “Probabilistic Evidentialist Reliabilism” to honor the three components that play a crucial role in it: evidence, reliability, and probability. One issue with relying on the notion of probability is the problem of logical omniscience. Consider the fact that one of the axioms in the usual Kolmogorov presentation of the probability calculus states that all tautologies receive maximal probability. This means that a subject will automatically be justified in being certain of any tautology whatsoever, even those too complicated for the subject to even parse. Different approaches to the problem of logical omniscience have been attempted, but I will not adjudicate between them here. I will here just note that my own opinion on this issue is that the problem of logical omniscience is serious, and it does show that epistemic justification simply cannot have a probabilistic structure. However, that does not mean that we cannot appeal to probabilities in our

¹¹See the preface to the second edition of Carnap (1950).

theories, but just that we must be careful in what it is that we are modeling with them. Of course logical omniscience is not rationally required, but if what we are interested in a particular context is not the epistemology of logic, then we can safely bracket issues having to do with logical omniscience and rely on a probabilistic structure for epistemic justification.¹² The reader can verify that none of the arguments that follow trade on so bracketing those issues.

6. What is Pr?

The theory we have arrived at is very similar to one recently proposed by Tang (2016b), which in turn, as noticed by Pettigrew (ms), is at least extensionally equivalent to a theory proposed by Dunn (2015) (at least when Dunn's theory is supplemented as we have done here to deal with the generality problem) as well as to Pettigrew's own theory. The theories do not incorporate the account of evidence as undefeated experience that I gave in the previous section, but on the other hand they do answer a question that we have so far not considered: namely, what is the probability function in whose terms we defined justified credence? All three of Tang, Pettigrew and Dunn answer that it is an objective, non-evidential probability which is in question. It is objective because it does not measure the actual degrees of belief of any particular agent. But it is not an evidential probability function, because it is not a measure of the degrees of belief it is rational to have. Or, more precisely, it will turn out to be an evidential probability function in the end (after all, the theories are all theories of justified credence), but only because the theories have it that the evidential probability function just coincides with a non-evidential one. The theories are trying to explain evidential probabilities, and they do so by appealing to a non-evidential type of probability.

Pettigrew is the most explicit about this. He says:

Given Jenann Ismael's distinction between single-case objective probabilities and general objective probabilities, our notion falls under the latter heading (Ismael (2011)). They are what Elliott Sober (2010) calls macro-probabilities; they are what David Albert (2000) finds in classical statistical mechanics. For Ismael, single-case probabilities tend to be unconditional, and assign values to particular token events, such as this particular die landing six on this particular roll. They are sometimes called chances, and they are the sort of probabilities we find in quantum mechanics. They are the objective probabilities that propensity accounts and best-system analyses aim to explicate. General probabilities, in contrast, tend to be conditional, and they take as arguments a pair of event types, such as dice of a certain sort landing six given that they are rolled. These are the sorts of probabilities that are found in statistical mechanics and evolutionary biology. They are what frequentist accounts attempt to explicate.

And he goes on to provide some examples:

¹²See Williamson (forthcoming) for a development of this kind of view.

Crucially, for our purpose, non-trivial general probabilities are possible even in deterministic worlds, whereas non-trivial chances are not—indeed, it is part of what it means for a world to be deterministic that the chance of an event at any time is either 0 or 1. Thus, in such a world, any particular roll of any particular die either is determined to come up six or is determined not to come up six. Nonetheless, it is still possible in such a world that the general objective probability of a die with certain general physical properties landing six given that it is rolled is 1. Similarly, while it is determined by the deterministic laws whether any particular egg in any human reproductive system will survive to reproductive age or will not, there is nonetheless a non-trivial probability that an egg will survive to reproductive age, given that it is a human egg — and this is a general probability. And while it is determined by the deterministic laws whether any particular block of ice in warm water will or will not melt, there is nonetheless a non-trivial (though very high) probability that a block of ice will melt given that it is in warm water—again, this is a general probability.

Given that this is the nature of the probability function Pr , however, it is quite clear that it will render all of the theories that appeal to it materially inadequate—they will all have clear counterexamples. Moreover, the counterexamples will not be easily brushed-off as marginal or somehow not terribly relevant—they strike at the heart of the theories, and show that they are, simply put, wrong.

BonJour's clairvoyant counterexample to Reliabilism was designed to show that reliability is not sufficient for justification. I argued (in Comesaña (2009)) that a kind of Reliabilism that incorporates evidentialist themes is not vulnerable to BonJour's counterexamples. However, I think now that counterexamples of the same kind can be given even against sophisticated reliabilist theories that incorporate evidentialist notions, provided that they appeal to the kind of objective probability that Pettigrew is alluding to. BonJour's counterexamples rested on the fact that blind reliability doesn't provide justification. Adding the notion of evidence to Reliabilism, I thought, provided a way out of the blindness. But measuring reliability in terms of general objective probabilities blinds Reliabilism once again, even when inoculated with evidence.

Suppose that you know by seeing that a certain leaf has shape S . As it happens, the general objective probability of its being the leaf of an oak tree given that it has shape S is n . You, however, are no botanist, and know nothing about oaks and their leaves. Even if, by chance, you happen to assign credence n to the proposition that the leaf is an oak leaf, that credence will in no way be justified. Or suppose that you know that a certain patient has symptoms S . As it happens, the general objective probability of a person's having disease D given that they exhibit symptoms S is m . You, however, are no physician, and know nothing about disease D and its symptoms. Even if, by chance, you happen to assign credence m to the proposition that the patient exhibiting symptoms S has disease D , that credence will in no way be justified. Moreover, those credences will not be justified even if your assignment of credences is done on the basis of the evidence in question. Suppose that you base your credence about the leaf coming from an oak on the fact that it has shape S , and you base your credence in the patient's having disease D on the fact that he

exhibits symptoms S . Far from this making your credence assignments rational, it just highlights the role that blind luck is playing in your credence assignments—a kind of luck incompatible with justification.

Now, it may be replied that if you have no idea about the connection between the leaf shape and its provenance, or between the symptoms and its causes, then you are not really basing your credence assignments on the evidence. At best, the evidence is playing a merely causal role in those credence assignments—but even advocates of causal accounts of basing need to grant that the existence of a causal relation is not *sufficient* for epistemic basing. Perhaps the subject also needs to somehow appreciate the connection between the evidence and the doxastic attitude it justifies.

But this response cannot help the Evidentialist Reliabilist who appeals to an objective type of probability. To begin with, it has a decidedly internalist flavor which no reliabilist worth the name (not even an evidentialist one) should be comfortable with. But, more seriously, if we buy the idea that basing requires an appreciation of the bearing of the evidence on the target doxastic attitude, not any old appreciation will do—it will have to be a *justified* appreciation. For instance, if you also just happen to think that n % of leaves with shape S belong to oaks, that will in no way make your credence assignment more justified. And if your appreciation of the bearing of the evidence has to itself be justified, then this launches a regress that ends either with some evidence giving you doxastic justification even absent appreciation of its bearing, or with some such appreciation being justified non-evidentially. But if the regress is resolved the first way that just leaves the theory open to the original objection. For now we will have a case where the subject assigns a credence n to a proposition p on the basis of some evidence E , and while it is true that $Pr(p|E) = n$, interpreting Pr as a general objective probability, the subject has no idea about the connection between E and p . On the other hand, if the regress is resolved the second way, by saying that justified appreciations of the bearing of evidence on doxastic attitudes may themselves be non-evidentially based, that is just incompatible with the theories, which have the consequence that all justified attitudes are evidence-based.

Cohen's new evil demon objection to reliabilism was designed to show that reliability is not necessary for justification. I argued (in Comesaña (2002)) for a solution to that objection based on a two-dimensional semantics for "reliable." But the theories of Dunn, Tang and Pettigrew do not incorporate that detail of my position. As such, they are vulnerable to Cohen's objection—or a similar one at least. For suppose that a subject lives in a counter-inductive environment. For instance, suppose that the environment puts evolutionary pressure on bird species to have different colors.¹³ In that case, the general objective probability of the hypothesis that all ravens are black diminishes as the number of observed black ravens grows. Nevertheless, it would be irrational for a subject to become more and more convinced that not all ravens are black the more black ravens she observes.

Behind these specific counterexamples lies a more fundamental problem with appealing to this kind of objective probability, and that is that it is a contingent kind of probability.

¹³Compare Titelbaum (forthcoming) on the "Hall of different-colored birds".

The value of a particular conditional probability of this kind depends on the contingent regularities that obtain in the world. As such, it will only be rational to match one's credences to those probabilities after one learns about the correlations. But learning is itself an epistemic achievement, and involves rational belief. Therefore, the rationality of the doxastic attitudes cannot be explained in terms of that kind of probability.

7. Evidential Probability

If Pr cannot be the kind of objective probability function that Dunn, Tang and Pettigrew take it to be, then what is it? One obvious answer (one which I hinted at in Comesaña (2009)) is: the evidential probability function. The evidential probability function determines two things: what credence it is rational to assign to different hypotheses in the absence of any evidence for or against them, and (via the conditional evidential probability function) what credences it is rational to assign to propositions given certain evidence. It is also an objective kind of probability, but not contingent. It is the kind of probability that Carnap (1950) tried (in vain) to define in purely syntactic terms. Since then, many philosophers have equated the failure of the Carnapian project with a refutation of the existence of such an evidential probability function. But, of course, the failure of the Carnapian project does nothing of the sort: it just shows that the difference between, say, “green” and “grue” is not syntactic.¹⁴

It is hard to find explicit formulations of that kind of skepticism in print, but Titelbaum (2010) is an exception. In that paper, Titelbaum explains and generalizes the problem that Goodman's grue example (Goodman (1979)) poses for the existence of the evidential probability function. Titelbaum's argument explicitly involves the rejection of the claim that the shape of the evidential probability function can be gleaned *a priori*. But if the argument of the previous section against appealing to contingent probability functions is on the right track, then the right probability function could only be gleaned *a priori*. Titelbaum would then apply Modus Tollens and conclude that, therefore, there is no such evidential probability function. Titelbaum ends his paper with a nod towards subjective Bayesianisms according to which there is no unique evidential probability function, but any probabilistically coherent credence distribution is acceptable. Not to put too fine a point on it, that kind of subjectivism allows for *any* reaction to the evidence to be rational. That is, in my case at least, literally unbelievable. But I agree with Titelbaum on the options here: either we embrace the existence of an evidential probability function which cannot be determined syntactically, or we go full subjectivists—that is to say, relativists—about epistemology. This kind of epistemological relativism is particularly vulnerable to the self-undermining objection that most relativisms must face. That is to say, if Titelbaum is right, then his argument for the correctness of his position should command rational assent only for those who share Titelbaum's priors. There is of course, vastly more to say about these issues, and I say a little more in the next section.

The theory we have arrived at, then, is the following. First, the evidence a subject has at t

¹⁴I speak of “the” evidential probability function, thus committing myself to the uniqueness thesis in epistemology. As far as I can tell, however, the issues discussed here do not depend on this thesis.

is provided the undefeated experiences the subject has. Second, the credences a subject is justified in assigning at t are those determined by $Pr(-|E)$, where Pr is the evidential probability function and E is the subject's evidence at t .

Let us now see how the appeal to evidential probabilities can answer the problems for reliabilism that we argued cannot be answered by appeal to an objective probability function. BonJour's style of counterexample focused on cases where reliability is allegedly not sufficient for justification. But if reliability is measured in terms of the evidential probability function, then this kind of counterexample is impossible. For, simply put, if a subject is not justified in believing a proposition p despite having evidence which makes p sufficiently likely according to a probability function Pr , then Pr is not the evidential probability function. Similarly, Cohen's counterexamples require a situation where a subject is justified in believing a proposition p even though the subject's evidence makes p unlikely according to Pr —but, again, that is only possible if Pr is not the evidential probability function. More generally, as I argued in the previous section, both BonJour's and Cohen's counterexamples to reliabilism rely on measuring reliability according to a function whose values are contingent. But the values delivered by the evidential probability function are not contingent.¹⁵

Now, in one respect, the fact that measuring reliability according to the evidential probability function relieves Reliabilism of its problems with counterexamples is, of course, a welcome feature. In another respect, however, it might seem that the appeal to the evidential probability function makes the reliabilist answer to alleged counterexamples *too easy*. Relatedly, there is a circularity worry here: we are trying to give an account of epistemic justification, and we end up appealing to the evidential probability function, which just encodes under what conditions doxastic attitudes are justified. Moreover, we are not giving an independent specification of the evidential probability function, but just appealing to it, whatever it is. How, therefore, is this progress? I come back to this question below. First, however, we need to reconsider the role of evidence in Evidentialist Reliabilism.

8. Evidence and Credences

As mentioned before, Conee and Feldman themselves have a mentalistic conception of evidence. The rough idea seems to be that certain special mental states—experiences paradigmatically, but perhaps also apparent memories and other non-factive mental states—“start the ball rolling,” in that they are the foundations of all justified belief. More specifically, mentalism has it that these mental states are themselves evidence, and they

¹⁵In Comesaña (2010a) I argued that the mere appeal to evidence could answer BonJour's counterexample. To be more precise, I granted that maybe BonJour's counterexamples did show that reliability is only necessary for justification, but I didn't comment on the fact that this just means that Reliabilism thus conceived was at best only a partial account of evidential fit. In that same paper I adopted my previous answer to Cohen's new evil demon problem presented in Comesaña (2002). In effect, my proposal there is one way to make contingent reliable connections into necessary ones. Given the necessity of evidential probabilities, this more roundabout solution is not necessary.

give rise to evidentially (but non-inferentially) justified beliefs, which in turn can inferentially justify further beliefs. In what follows I talk of experiences exclusively, but I mean to leave it open that other mental states can also provide basic justification.

My worry about such a conception of evidence is that it is simply not plausible that the fact that a subject is undergoing a certain mental state is evidence for propositions about things completely unrelated to that subject and her mental states. Suppose that you know nothing about a certain subject, and I tell you that she is undergoing a certain mental state and ask you whether that is evidence (for her, presumably, but there is here an interesting question about the privacy of evidence under the mentalist construal) for the proposition that there is a snowball in a certain spatiotemporal position (right in front of the subject, say). Knowing nothing of the subject and the nature of her mental states, the experience is completely irrelevant to the existence of snowballs. It is just as if I told you that there is a certain sentence written in the sand, and asked you whether that sentence is evidence that there is life on Mars. The obvious answer is that the sentence is evidentially irrelevant to the existence of life on Mars. Now, I am leaving aside something which many philosophers will find deeply relevant to our question, which is the fact that experiences have *content*. Some philosophers, of course, will deny this—but I don't. But even if I tell you that the experience that the subject is undergoing has a certain content, and even if I tell you that the content just is that there is a snowball right in front of her, you should not take that to be evidentially relevant to whether there is indeed a snowball in front of her. After all, even if I told you that the sentence written in the sand is "There is life on Mars," you wouldn't take that to be evidentially relevant to whether there is life on Mars. Of course, if I told you that someone who knows a lot about the issue wrote that sentence down with the intention to inform people, then that would of course be evidence that there is life on Mars. Similarly, if I told you more about this subject, perhaps for instance that she belongs to a species that has evolved in and is well adapted to her environment, or perhaps if I just told you that her experiences are very reliable, then that might indeed be good evidence that there is a snowball in front of her. But the resulting theory of evidence is not mentalism, but rather something closer to classical foundationalism.¹⁶ The resulting theory has it that a pair of propositions, namely that a subject has an experience with the content that *p* and that the subject's experiences are reliable, are evidence for *p*. The mentalist account has it that the experience itself, not propositions about it, are the evidence.¹⁷

An alternative is that mental states like experience can provide, but do not consist of, evidence. One version of this alternative adopts Williamson's equation of evidence with knowledge and has it that the experience provides the subject with the proposition that is its content as evidence when the subject knows the proposition on the basis of that experience. Another version, which I prefer, has it that truth is irrelevant to whether the experience provides the subject with evidence, and that it does so as long as it provides her with justified belief. Call this account of evidence, "Propositionalism."¹⁸ Now,

¹⁶Which is not to say that some self-avowed evidentialist would not be happy with a theory like this—see, for instance, McCain (forthcoming).

¹⁷The argument in this paragraph is developed in more detail in Comesaña (2015).

¹⁸For more on this, see Comesaña and McGrath (2014), Comesaña (2015), Comesaña and McGrath (2016), and the article by McGrath in this volume.

this might seem to be viciously circular. After all, we started out with the objective of giving an account of justified belief, and we ended up with a theory that bottoms out in evidence which is provided by experiences when they give rise to justified belief. But if there is any circularity here, it is of a kind that every theory under consideration here shares. Consider first the mentalist conception of evidence. It holds that experience provides *prima facie* justification for a belief in its content, which turns into all things considered justification when undefeated. Williamson’s view has it that an experience provides the proposition that is its content as evidence when it provides knowledge—and that would happen, of course, only if the experience provides justification for belief in its content. So both mentalism and $E = K$ have it that experiences provide or consist of evidence when they justify belief in their contents. We are all, therefore, in the same boat when it comes to the project of specifying why and under what conditions experiences provide justified belief in their contents, and none of the three views should be judged to be viciously circular for holding that the evidential chain begins when that happens. Every view, then, needs some kind of account of when an experience with the content that p results in a justified belief that p . I use “results” as a term that is neutral between the mentalist account of evidence—according to which the experience itself is evidence for the belief—and the Williamsonian and Propositionalist accounts, according to which experience provides us with, but does not consist of, evidence. Let us say that, when an experience results in a justified belief, the experience is undefeated. Everyone, then, needs an account of when experience is undefeated. Fortunately, we need not wait to have such a complete account before we use the notion of undefeated experience in our theorizing, for we understand perfectly well what that notion means. Now, maybe Reliabilists have a legitimate complaint here, because they would say that an experience provides the subject with a justified belief in its content just in case there is a reliable connection between the subject’s having the experience and the truth of its content. But this is at best a Reliabilist account of undefeated experience—and likely not a correct one. Goldman (1979) himself provided a reliabilist account of defeat in general, but it is not at all clear that he succeeded.

The theory we have arrived at is the following: experiences provide the subject with an initial corpus of propositions as evidence, and they do so provided that they justify the subject in believing their contents; this initial corpus of evidence can then justify belief in a further proposition when the conditional probability of that proposition given a suitable subset of the initial corpus is high enough. A natural thought here would be to think that propositions justified downstream of experience can then join forces with the basically justified propositions and justify still more propositions, the edifice of justified propositions growing under its own steam, so to speak. But although this picture is a very traditional one, I do not think that it withstands scrutiny.¹⁹

Suppose that a subject possesses some corpus of propositions E as basic evidence (at a given time). Suppose, in addition, that $1 > Pr(p|E) \geq r$, so that p is justified for S by E but not to the maximal degree. Finally, suppose also that although $Pr(q|p) \geq r$, $Pr(q|E) < r$, so that whereas q is justified for S if p is part of her evidence, it is not

¹⁹Sosa (2016) criticized Evidentialist Reliabilism precisely on the basis that it, together with Evidentialism, assumed that all beliefs are evidentially justified. As the theory of evidence deployed here and developed further in the articles cited in the previous footnote show, I agree with Sosa.

if it is merely E that is part of S 's evidence. In that case, the rational credence for S to have in q is clearly $Pr(q|E)$ —to assign it $Pr(q|p)$ would be irrational overconfidence.²⁰ Moreover, this irrational overconfidence in q can lead to irrational actions. Therefore, we shouldn't treat propositions justified by our basic evidence as themselves being part of our evidence. Our evidence, therefore, consists of those propositions which are justified for us, but not evidentially justified.

The resulting theory can be formulated as follows:²¹

Coarse-Grained Evidentialist Reliabilism: A belief that p by S is justified if and only if:

Either:

1. S 's undefeated experiences provide him with p ; or
- 2a. S 's undefeated experiences provide him with E ;
- 2b. the belief that p by S is based on E ;
- 2c. $Pr(p|E) \geq r$.
- 2d. There is no more inclusive body of evidence E' had by S such that $Pr(p|E) < r$.

I call the resulting theory “Coarse-Grained Evidentialist Reliabilism” because it accounts only for the rationality of coarse-grained doxastic attitudes, and not of fine-grained ones. In that respect, it is inferior to plain old Evidentialism, which applied to all doxastic attitudes. There is a very natural way, however, to transform Coarse-Grained Evidentialist Reliabilism into a theory that applies to credences:

Fine-Grained Evidentialist Reliabilism: A credence x in p by S is justified if and only if:

Either:

1. S 's undefeated experiences provide him with p ; or
- 2a. S 's undefeated experiences provide him with E ;
- 2b. S 's credence x in p is based on E ;
- 2c. $Pr(p|E) = x$.
- 2d. There is no more inclusive body of evidence E' had by S such that $Pr(p|E) \neq x$.

What of the complain, raised before, that the theory conflates justification as firmness with justification as increase in firmness? The complaint doesn't apply to Fine-Grained

²⁰The issues here are intimately related to the “easy knowledge” problem—see Cohen (2002).

²¹The definitions are implicitly relativized to a time. That doesn't mean that the theory is a version of “time-slice” epistemology, according to which which doxastic attitudes are justified at a time supervenes on the subject's mental states at a time, for it leaves it open that past experiences may provide subjects with present evidence.

Evidentialist Reliabilism because it does not deal with a threshold notion of justification. What we gained by having a theory that applies to credences, however, we lost in that it no longer applies to full beliefs. Obviously, however, we can join both theories just by adding that a justified credence x in p counts as justified belief in p if and only if $x \geq r$. This makes it even more explicit than it was before that the conception of justified belief implicit in Coarse-Grained Evidentialist Reliabilism follows the Lockean thesis according to which justified belief just is justified credence above a threshold. An alternative is to say that the only propositions the subject is justified in fully believing (and not just in assigning some high credence) are those that are part of her evidence. This also has obvious costs, but this is not the place to deal with them.²²

9. Conclusion

Is Fine-Grained Evidentialist Reliabilism well named? It is in the sense of honoring its ancestry, but the resulting view is, in important respects, neither Evidentialist nor Reliabilist. (It is also, of course, more than a mouthful.)

Starting with Evidentialism, the most fundamental difference between the views is that Fine-Grained Evidentialist Reliabilism embraces the possibility of non-evidentially justified beliefs.²³ Relatedly, Fine-Grained Evidentialist Reliabilism has it that a subject's evidence is constituted exclusively by those propositions provided as evidence by his undefeated experiences (remember that we are treating "experiences" as somewhat of a placeholder for all non-factive mental states which can provide evidence). For the Evidentialist, remember, there are two fundamentally different kinds of evidence: evidence can consist of propositions, which are had as evidence only if justifiedly believed, or of experiences, for which there is no distinction between evidence and its possession. Fine-Grained Evidentialist Reliabilism, on the other hand, has a unified conception of evidence and its possession. Conversely, whereas for the Evidentialist all propositions are evidentially justified, for the Fine-Grained Evidentialist Reliabilist some propositions are justified but not on the basis of any evidence. Both kinds of theory therefore posit some kind of bifurcation: one in the notion of evidence and its possession, the other on the ways beliefs can be justified. I have briefly alluded at the arguments for preferring the second kind of bifurcation. In his contribution to this volume, Matt McGrath concentrates on precisely this issue, and mounts an argument against Evidentialism on it.

Fine-Grained Evidentialist Reliabilism replaces the Evidentialist notion of fit with an appeal to the evidential probability function. So, in this respect, while Fine-Grained Evidentialist Reliabilism might not be better off than Evidentialism, it clearly isn't worse off. I conceived of the Reliabilist part of Evidentialist Reliabilism as providing an answer

²²Christensen (2004) argues that, given those problems, we should just abandon coarse-grained epistemology in favor of fine-grained epistemology. Pragmatic encroachment à la Fantl and McGrath (2002) might help deal with some of the problems, but it is of course itself a very controversial theory.

²³That said, if an Evidentialist is happy with saying that the item of evidence is the content of the experience, and that it is possessed in virtue of the subject's undergoing the experience, then so be it. I have no problem with calling my view "Evidentialist", but I want to make clear what conception of evidence and its possession I am arguing for.

to the question of fit, but I didn't pay sufficient attention to the question of how to measure reliability. If we measure reliability by an objective probability function, then the old counterexamples to Reliabilism come back with a vengeance, and the admixture of Evidentialism will not help. Therefore, while the versions of Reliabilism advocated by Dunn, Tang and Pettigrew do indeed provide us with a non-circular account of justification, that account is simply materially inadequate. Moreover, its material inadequacy can be traced back precisely to the fact that they conceive of Pr as a contingent function, and as such the only justified way to match our credences to it is by learning about those contingent correlations. But part of what we need to explain when we explain epistemic justification is precisely how it is that we are justified in learning about those correlations.

An explicit part of Goldman's project in epistemology is to provide a *reductivist* account of epistemic justification. Insofar as the kind of Evidentialist Reliabilism defended here appeals to the notion of evidential probability, it cannot fulfill that promise. As I see it, however, reductivism of that kind comes at the cost of material inadequacy.

Fine-Grained Evidentialist Reliabilism is a theory with its own theoretical commitments, and, I think, well worth exploring.²⁴

References

- Albert, David. 2000. *Time and Chance*. Harvard University Press.
- Alston, William. 1988. "An Internalist Externalism." *Synthese* 74: 265–83.
- BonJour, Lawrence. 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* V: 53–73.
- Carnap, Rudolph. 1950. *Logical Foundations of Probability*. University of Chicago Press.
- Christensen, David. 2004. *Putting Logic in Its Place*. Oxford University Press.
- Cohen, Stewart. 1984. "Justification and Truth." *Philosophical Studies* 46: 279–95.
- . 2002. "Basic Knowledge and the Problem of Easy Knowledge." *Philosophy and Phenomenological Research* 65 (2): 309–29.
- Cohen, Stewart, and Juan Comesaña. 2013. "Williamson on Gettier Cases and Epistemic Logic." *Inquiry* 56 (1): 15–29.
- . forthcoming. "Rationality and Truth." In *The New Evil Demon*. Oxford University Press.
- Comesaña, Juan. 2002. "The Diagonal and the Demon." *Philosophical Studies* 110: 249–66.
- . 2006. "A Well-Founded Solution to the Generality Problem." *Philosophical Studies* 129 (127–47).

²⁴Many thanks to Kevin McCain for helpful comments on a draft of this paper.

- . 2009. “What Lottery Problem for Reliabilism?” *Pacific Philosophical Quarterly* 90 (1): 1–20.
- . 2010a. “Evidentialist Reliabilism.” *Noûs* 44 (4): 571–600.
- . 2010b. “Reliabilism.” In *The Routledge Companion to Epistemology*, edited by Sven Bernecker and Duncan Pritchard. Routledge.
- . 2015. “Falsehood and Entailment.” *Philosophical Perspectives* 29 (1): 82–94.
- . 2017. “On Sharon and Spectre’s Argument Against Closure.” *Philosophical Studies* 174 (4): 1039–46.
- Comesaña, Juan, and Matthew McGrath. 2014. “Having False Reasons.” In *Epistemic Norms*, edited by Clayton Littlejohn and John Turri. Oxford University Press.
- . 2016. “Perceptual Reasons.” *Philosophical Studies* 173 (4): 991–1006.
- Conee, Earl, and Richard Feldman. 1985. “Evidentialism.” *Philosophical Studies* 48.
- . 1998. “The Generality Problem for Reliabilism.” *Philosophical Studies*, no. 89: 1–29.
- . 2008. “Evidence.” In *Epistemology: New Essays*, edited by Quentin Smith, 83–104. Oxford University Press.
- Dunn, Jeff. 2015. “Reliability for Degrees of Belief.” *Philosophical Studies* 172 (7): 1929–52.
- Fantl, Jeremy, and Matthew McGrath. 2002. “Evidence, Pragmatics, and Justification.” *Philosophical Review* 112 (1): 47–67.
- Goldman, Alvin. 1979. “What Is Justified Belief?” In *Justification and Knowledge*, edited by George Pappas, 1–23. Dordrecht: Reidel.
- . 2011. “Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism’s Troubles, Reliabilism’s Rescue Package.” In *Evidentialism and Its Discontents*, edited by Trent Dougherty. Oxford University Press.
- Goodman, Nelson. 1979. *Fact, Fiction and Forecast*. Harvard University Press.
- Ismael, Jennan. 2011. “A Modest Proposal About Chance.” *Journal of Philosophy* 108 (8): 416–42.
- McCain, Kevin. forthcoming. “Explanationist Aid for Phenomenal Conservatism.” *Synthese*.
- Pettigrew, Richard. ms. “What Is Justified Credence?”
- Sober, Elliott. 2010. “Evolutionary Theory and the Reality of Macro Probabilities.” In *The Place of Probability in Science*, edited by Ellery Eells and J. Fetzer, 133–61. Springer.
- Sosa, Ernest. 2016. “Process Reliabilism and Virtue Epistemology.” In *Goldman and His Critics*, edited by Brian McLaughlin and Hilary Kornblith. Wiley-Blackwell.

- Tang, Weng Hong. 2016a. "Reliabilism and the Suspension of Belief." *Australasian Journal of Philosophy* 94 (2): 362–77.
- . 2016b. "Reliability Theories of Justified Credence." *Mind* 125 (497): 63–94.
- Titelbaum, Michael. 2010. "Not Enough There There* Evidence, Reasons, and Language Independence." *Philosophical Perspectives* 24: 477–528.
- . forthcoming. *Fundamentals of Bayesian Epistemology*. Oxford University Press.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.
- . 2013. "Gettier Cases in Epistemic Logic." *Inquiry* 56 (1): 1–14.
- . forthcoming. "Justification, Excuses and Sceptical Scenarios." In *The New Evil Demon*, edited by Julien Dutant and Fabian Dorsch. Oxford University Press.
- . forthcoming. "Model-Building in Philosophy." In *Philosophy's Future: The Problem of Philosophical Progress*, edited by Russell Blackford and Damien Broderick. Blackwell.