Justification, Normalcy and Evidential Probability

Martin Smith

My concern in this paper is with a certain, pervasive picture of epistemic justification. On this picture, acquiring justification for believing something is essentially a matter of minimising one's risk of error – so one is justified in believing something just in case it is sufficiently likely, given one's evidence, to be true. This view is motivated by an admittedly natural thought: If we want to be fallibilists about justification then we shouldn't demand that something be *certain* – that we *completely* eliminate error risk – before we can be justified in believing it. But if justification does not require the complete elimination of error risk, then what could it possibly require if not its minimisation? If justification does not require epistemic certainty then what could it possibly require if not epistemic likelihood? When all is said and done, I'm not sure that I can offer satisfactory answers to these questions – but I will attempt to trace out some possible answers here. The alternative picture that I'll outline makes use of a notion of *normalcy* that I take to be irreducible to notions of statistical frequency or predominance.

NOTE: This paper is a reworking of some aspects of a previous paper of mine – 'What else justification could be' published in *Noûs* in 2010. I'm currently in the process of writing a book developing and defending some of the ideas from this paper. What follows will, I hope, fall into place as one of the chapters of this book – though it is still very much at the draft stage. Comments are welcome.

I. INTRODUCTION

Some philosophers have claimed that, alongside standard Gettier cases, lottery cases provide further, vivid counterexamples to the traditional analysis of knowledge as justified, true belief (see Hawthorne, 2003, pp9, Pritchard, 2007, pp4). They reason along the following lines: Suppose that I hold a single ticket in a fair lottery of one million tickets. Suppose that I am convinced, purely on the basis of the odds involved, that my ticket won't win. Do I *know* that my ticket won't win? Intuitively, I don't know any such thing, even if it happens to be true. Presumably, though, I have plenty of *justification for believing* that my ticket won't win – after all, given my evidence, this proposition has a 99.9999% chance of being true. How much more justification could one want? If I'm not justified in believing that my ticket won't win, then surely none of us are justified in believing much at all. Here is a case, then, in which a justified, true belief fails to qualify as knowledge.

This argument seems straightforward enough, and yet there are reasons for being a little uneasy. On reflection, lottery cases seem somehow *different* from standard Gettier cases. Consider the following for comparison: I wander into a room, undergo a visual experience as of a red wall and come to believe that the wall is red. In actual fact the wall *is* red but, unbeknownst to me, it is bathed in strong red light emanating from a hidden source, such that it would have looked exactly the same even if it had been white. Intuitively, I do not know, in this case, that the wall is red, in spite of the fact that my belief is both justified and true.

We can observe a number of apparent differences between these two cases. In particular, while my belief in the Gettier case fails to actually qualify as knowledge, it nevertheless seems to be a good or promising *candidate* for knowledge – and would have been knowledge if only conditions in the world had been more obliging. My belief in the lottery case, however, doesn't seem to be the sort of belief that could *ever* qualify as knowledge. In the Gettier case, the problem seems to lie with the *world* – and funny, abnormal goings on therein. In the lottery case, the problem seems to lie with *me* and the way in which I form my belief (see Ryan, 1996, pp136, 137).

The Hawthorne/Pritchard argument betrays a commitment to a certain, quite pervasive, way of thinking about epistemic justification. The picture is something like this: For any proposition P we can always ask how *likely* it is that P is true, given present evidence. The more likely it is that P is true, the more justification one has for believing that it is. The less likely it is that P is true, the less justification one has for believing that it is. One has justification simpliciter for believing P when the likelihood of P is sufficiently high and the risk of ~P is correspondingly low. Call this the *risk minimisation* conception of justification.

So long as we're relatively liberal in our interpretation of 'evidence' and 'likelihood', this general picture can be identified in the work of a very broad range of epistemologists. Sometimes it is explicitly set-out (see Russell, 1948, chap. VI, Chisholm, 1957, pp28, Goldman, 1986, section 5.5, Fumerton, 1995, Pryor, 2004, pp350-351, 2005, BonJour, 2010). More often, it is left largely implicit, as in the above reasoning. I don't know of any detailed arguments in favour of this picture, but the following thought is admittedly quite compelling: Most epistemologists are *fallibilists* of one kind or another and think that a belief can be justified even if we don't have evidence that makes it certain – even if we haven't completely eliminated all risk of error. But if justification can fall short of evidential certainty, then what else could it possibly *be* if not evidential likelihood? If justification does not require the complete elimination of error-risk, then what else could it possibly require if not its minimisation?

It's tempting to think, when in this mindset, that the risk minimisation conception is not really a substantial claim at all – rather it just serves to *define* a basic fallibilist notion of justification that precedes any substantial epistemic theorising. But whatever else one might think of this picture it can't really be *trivially* true. The idea that beliefs can be appraised as permissible or impermissible, in a distinctively epistemic sense, is one that is betrayed by a broad range of ordinary practices; by our evaluating the adequacy of methods of inquiry and in determining when inquiry into a matter might reasonably cease, by our condemning assertions as ungrounded or premature and in our criticising actions for the beliefs upon which they are based. If there is a 'basic' notion of justification it is just this: One has justification for believing a proposition P iff one is epistemically permitted to believe that P. One has more justification for believing P than Q iff one is more strongly epistemically permitted to believe that P than to believe that Q. Even if it is true that the minimisation of error risk is what makes belief epistemically permissible, it is not so by stipulation.

In this paper, I shall offer an argument against the risk minimisation conception of justification. I shall also begin the task of outlining an alternative way in which we might think about justification – *another thing* that justification could be, if not risk minimisation. My alternative makes use of a notion of *normalcy* that is irreducible to notions of statistical frequency or predominance.

I assume here that evidence is propositional – that one's body of evidence consists of a stock of propositions (Williamson, 2000, section 9.5). I won't be defending any particular account of when a proposition qualifies as part of one's body of evidence – indeed, everything that I say here will be compatible with a number of different ways of thinking about this. In particular, everything that I say here can be reconciled with generous accounts of evidence, such as Williamson's knowledge account (Williamson, 2000, chap. 9), and with more austere psychologistic accounts (Lewis, 1996), as well as a range of further views.

Even if we fix upon a particular account of evidence, there is still room for substantial disagreement over the nature of evidential probability – over what it takes for a piece of evidence to confer a particular probability value upon a proposition. *Externalists* about evidential probability hold that evidential probability values are, in general, the product of contingent facts – and in particular, of facts about *relative frequencies*. On this view, roughly speaking, the probability of a proposition P, given evidence E, will be determined by the frequency with which the truth of E-type propositions is accompanied by the truth of P-type propositions in actual and close counterfactual circumstances (see Russell, 1948, chap. VI Alston, 1988, section I). *Internalists*, on the other hand, conceive of evidential probability values as reflecting necessary, internal connections between propositions (see Keynes, 1921, Fumerton, 1995, 2004, Conee and Feldman, 2008, section 1.5). If E confers a certain probability value upon P then, according to the internalist, it could not have failed to do so.

I won't take for granted any particular account of evidential probability. Neither will I assume that propositions always have precise, well-defined evidential probability values relative to any possible body of evidence (White, 2005, Williamson, 2000, pp211)[1]. Our intuitions about evidential probability undoubtedly seem clearest when our evidence is of a broadly 'statistical' nature – when it concerns objective chances or proportions or frequencies or the properties of some putatively random process. I shall largely limit myself to clear-cut

[1] Leplin (2009, chap. 6) argues that most of the ordinary propositions that we believe have *no probabilities whatsoever*. The claim can be interpreted in a variety of different ways – but any suggestion that the evidential probabilities of most ordinary propositions are literally *undefined* is, I think, very difficult to accept. Suppose I'm sitting in my office with the curtains closed listening to the howling wind and what sound like water droplets striking my window pane. It is, of course, difficult in such circumstances to assign any principled, precise evidential probability value to the proposition that it's raining. But if this value really were undefined, then it would be impossible even to *compare* the evidential probability of this proposition to that of others – and yet this clearly seems possible. Surely, given my evidence, it's *more* likely that it's raining than sunny, for instance. My objections to the risk minimisation picture, in any case, will not rest upon any general doubts about the notion of evidential probability or its scope. I'll have more to say about these issues in section IV.

cases such as these. I will have more to say, both about evidence and evidential probability, in section IV.

I am taking it for granted here that believing a proposition involves a commitment over and above merely regarding it as likely. In one sense this seems obvious, and there is little temptation to think otherwise – to claim that P is *likely* is precisely to avoid committing to its truth[2]. I can regard it as *extremely* likely that, say, the number of stars in the universe is a composite number or that ticket #542 will lose the lottery without actually believing either of these things. In his *Rules for the Direction of the Mind*, Descartes famously advised that we should never believe that which is merely probable (rule II). I take it for granted that such advice is, at the very least, *intelligible*. (In fact, I am inclined to think that, when appropriately construed, Descartes' advice is perfectly *sound* – more on this later). If the evidential probability of proposition P given evidence E is high, then it is, in some sense, appropriate for any subject in possession of this evidence to offer a generous estimate of P's probability. Whether it's permissible for a subject, so situated, to believe that P is true is a further question – and one that may, as I hope to show, have different answers depending upon the character of E.

[2] There are, of course, difficult questions about the relationship between outright belief and degrees of belief – and these are not questions that I mean to prejudge here. According to one view, sometimes dubbed the 'Lockean Thesis', outright belief corresponds to a high degree of belief – one counts as believing a proposition just when the degree of belief that one invests in it is suitably high (see, for instance, Foley, 1992, 2009). I don't mean to assume here that such a view is incorrect – only that it cannot be combined with a view on which degrees of belief are equated with overt probability estimates.

II. PROBLEMS FOR RISK MINIMISATION

There are two simple theses behind what I've been calling the *risk minimisation* conception of justification. One of these concerns categorical attributions of justification and the other concerns comparisons:

(a) One has justification for believing a proposition P iff P is likely, given one's evidence, to be true.

(b) One has more justification for believing a proposition P than a proposition Q iff P is more likely, given one's evidence, than Q is.

A further claim is needed in order to make these intelligible: A body of evidence furnishes all, or most, propositions with evidential probabilities – probabilities that are usually taken to be describable by a classical (Kolmogorovian) probability function[3]. This is, perhaps, more usefully thought of as a precondition for the intelligibility of theses (a) and (b) rather than a further thesis alongside them. Talk of a proposition being 'likely' in thesis (a) should be further clarified in terms of a proposition having an evidential probability that exceeds some threshold t that may be variable and/or vague – but the details of this will matter little for present purposes.

---

[3] A probability function is defined over a set of propositions that includes a 'maximal' proposition, entailed by each of the others in the set, and is closed under negation and disjunction. Propositions are standardly modelled, for this purpose, as subsets of a set of possible worlds W, with W itself playing the role of the maximal proposition. A classical probability function Pr is any function that assigns real numbers to the propositions in this set and satisfies Kolmogorov's three axioms:

(i) $Pr(P) \geq 0$

(ii) $Pr(W) = 1$

(iii) $Pr(P \vee Q) = Pr(P) + Pr(Q)$, for any P and Q such that P and Q are inconsistent.

If the domain of the function taken to be infinite, and closed under countable disjunction then a fourth axiom is sometimes added:

(iv) If P_i, i = 1, 2 ... is a sequence of pairwise inconsistent propositions then $Pr(\vee_i P_i) = Pr(P_1) + Pr(P_2) + Pr(P_3)$ ...

A probability function that satisfies (iv) is known as *countably additive*.

It isn't an essential part of the risk minimisation conception, as I see it, that evidential probabilities be Kolmogorovian (though they are standardly taken to be). The view, as I understand it, can accommodate some difference of opinion regarding the formal behaviour of evidential probabilities. I will mention another sort of approach in the final section.

One possible argument against the risk minimisation conception proceeds from the principle of *multiple premise closure*. It's very plausible that deduction, of all things, should be an epistemically secure way of expanding one's set of beliefs – that deduction should never take us from the realm of justified belief into the realm of pure conjecture. One way of attempting to make this rough intuition precise is the following: If one has justification for believing each of a set of propositions, then one has justification for believing any proposition that deductively follows from them.

Multiple premise closure is in clear tension with thesis (a) of the risk minimisation conception. Multiple premise deductions can aggregate risk – the risk of error to which I expose myself in believing the conclusion of a multiple premise deductive inference may be higher than the risk of error to which I expose myself in believing any of the premises, taken individually. That is, the evidential probability of the conclusion of a multiple premise deduction can be lower than the evidential probabilities of each individual premise and, as such, the evidential probability of the conclusion can dip below the threshold for justification, even if the evidential probability of each premise surpasses it.

If multiple premise closure fails, then there must be possible situations in which it is epistemically permissible to believe the premises of a deductive argument and to accept its validity and yet epistemically *im*permissible to believe its conclusion. As curious as this seems, it won't be a part of my case against the risk minimisation conception. Indeed, I shall remain officially neutral on the principle of multiple premise closure here.

My argument begins, rather, from an example, adapted from one originally devised by Dana Nelkin (2000, pp388-389): Suppose that I have set up my laptop such that, whenever I turn it on, the colour of the background is determined by a random number generator. For one value out of one million possible values the background will be red. For the remaining 999 999 values, the background will be blue. One day I arrive at a desk in the library and turn on my laptop. Moments before the background appears, I spot my friend Bruce at a nearby desk and wander over to say hello.

Bruce is already working away on his laptop and, when I arrive, I immediately *see* that it's displaying a blue background and come to believe that it is. Let's suppose, for the time being, that my relevant evidence consists of two propositions (I will consider other ways of describing the evidential situation in due course):

(E_1) It visually appears to me that Bruce's laptop is displaying a blue background.

(E_2) It is 99.9999% likely that my laptop is displaying a blue background.

Here are a few preliminary observations about this case: If I were to believe that my laptop is displaying a blue background before returning to my desk, it would be natural to describe this belief as a *presumption* (albeit a very safe one), while it does not seem at all natural to describe my belief about Bruce's laptop in these terms. Second, my belief about Bruce's laptop would appear to be a very promising candidate for knowledge – indeed, it *will* be knowledge, provided we fill in the remaining details of the example in the most natural way. If I were to believe that my laptop is displaying a blue background, this belief would never constitute knowledge even if it happened to be true. If my battery died before I got back to my desk, I might well think to myself 'I guess I'll never know what colour the background really was'. But if Bruce's battery died I certainly wouldn't think this about the background colour on his laptop.

If someone were to ask me 'What colour is the background on Bruce's laptop?', I would be perfectly epistemically entitled to reply 'It's blue'. But if someone were to ask me the same question about *my* laptop, it seems as though I ought to be more circumspect, and say something along the lines of 'It's overwhelmingly likely that the background is blue – but I haven't actually seen it'. Presumably, this is what I ought to *believe* too. I'm not required to do any further investigation into the background colour displayed by Bruce's laptop – even though I easily *could* by, for instance, asking others to have a look. But I ought to do more investigation into the background colour displayed by my laptop – by, for instance, going and having a look myself – before I rest on my laurels.

The implication of these considerations seems clear enough: I have justification for believing that Bruce's laptop is displaying a blue background, but I don't have justification for believing that my laptop is displaying a blue background. In spite of this, the proposition that my laptop is displaying a blue background is more likely, given my evidence, than the proposition that Bruce's is. While E_1 does make it highly likely that Bruce's laptop is displaying a blue background, it clearly does not *guarantee* that it is. After all, I could be hallucinating, or I could have been struck by colour blindness, or I could be subject to some strange optical illusion etc. No doubt these are all rather unlikely – but presumably the likelihood, given evidence E_1, that Bruce's laptop is displaying a blue background would be no where near as high as 99.9999%. This, of course, is precisely how likely it is that my laptop is displaying a blue background, given evidence E_2[4]. In believing that Bruce's laptop is displaying a blue background, I am actually running a *higher* risk of error than I would be in believing the same thing about my laptop.

The judgment that I lack justification for believing that my laptop is displaying a blue background is, on its own, in tension with thesis (a) of the risk minimisation conception. This proposition is clearly *very likely*, given my evidence. One could, perhaps, try to preserve (a) by arguing that the likelihood threshold t should, for whatever reason, be set very high in the case that I've described. Bringing into play the judgment that I *have* justification for believing that Bruce's laptop is displaying a blue background effectively blocks this kind of manoeuvre. We cannot find a likelihood threshold such that it's likely that Bruce's background is blue, unlikely that my background is blue and yet *more likely* that my background is blue than Bruce's.

When it comes to thesis (b) of the risk minimisation conception, we do not, in fact, even need the judgment that I have justification for believing that the background on Bruce's laptop is blue, or the judgment that I lack justification for believing that the background on my laptop is blue. All that we need is a weaker *comparative* judgment – that I have *more* or better justification for believing that Bruce's background is blue than I do for believing that

---

[4] One could think of this as an instance of David Lewis's 'principal principle' (Lewis, 1980). Suppose my only evidence relevant to a proposition P is that the objective probability of P at time n = x. Suppose, in particular, that I have no relevant evidence pertaining to things that happened after n. According to the principal principle, roughly speaking, the evidential probability of P, for me, will also be equal to x. Time n, in this case, could be thought of as the time at which I turn on my laptop.

Whether my evidence, in the case described, should be interpreted, ultimately, as evidence pertaining to an *objective* probability value is, of course, debatable. I needn't take any stand on this – all that is important, for present purposes, is that its probabilistic bearing be clear-cut. It may be that it is best interpreted, at the end of the day, as evidence about a propensity or an expected frequency or some such. It's worth noting though that if we *ever* have evidence regarding objective probability values, then cases like the one described must be amongst the clearest cases in which we do.

mine is. This comparative judgment, then, is strictly enough to refute the letter of the risk minimisation conception.

The laptop example prompts a cluster of intuitions that don't seem to fit in well with the risk minimisation picture. It may be rather tempting, however, for one to simply disregard such intuitions as confused or naïve. Perhaps we are simply accustomed to relying upon perception in such matters and suspending any scruples about its fallibility. Once we do reflect carefully upon the fallibility of perception, so this thought goes, these troublesome intuitions are exposed as a kind of groundless prejudice. I'm not entirely convinced that this is the wrong thing to say – but I do strongly suspect that it is.

It may help to consider a further, somewhat less contrived, example discussed by Enoch, Fisher and Spectre (forthcoming) and Redmayne (2008), amongst others: A bus causes some harm on a city street – it damages a car or injures a pedestrian or some such. In the first scenario, an eye witness testifies that the bus was owned by the Blue-Bus company. In the second scenario, there is no eye-witness, but there is some unusually strong statistical evidence regarding the distribution of buses in the relevant area – evidence to the effect that 95% of the buses operating in the area, on the day in question, were owned by the Blue-Bus company.

Testimony, as we all know, is not perfectly reliable – particularly when it comes to testimony concerning an event of this kind. The eye witness in scenario one could have been mistaken about what she saw or could, indeed, have deliberately concocted a lie in order to smear the Blue-Bus company. These possibilities may not be likely, but the probability that it really was a Blue-Bus bus that was involved, given the eye witness testimony, is plausibly not as high as 95% – and this, presumably, is exactly how probable the proposition is, given the strong statistical evidence available in scenario two[5].

In spite of this, so long as we don't have any positive reason to think that the eye witness in scenario one is mistaken or lying, it would be perfectly reasonable for us to take this testimony at face value and conclude that it was a Blue-Bus bus that was involved. Indeed, if this is our only relevant evidence, it would seem reasonable for us to assert and to act upon this conclusion. In scenario two, on the other hand, all that we seem entitled to conclude is that it was *very likely* to have been a Blue-Bus bus that was involved. To conclude that it *was* a Blue-Bus bus and to assert or act upon this proposition would seem premature (see Enoch, Fisher and Spectre, forthcoming, Redmayne, 2008).

This is not just a philosopher's intuition – and neither is it purely academic. Under prevailing legal practice, the eye witness testimony available in scenario one could serve as a legitimate basis for a positive finding to the effect that the bus involved was a Blue-Bus bus – and an associated finding of liability. In contrast, it would not be legitimate to base a finding of liability upon the kind of statistical evidence available in scenario two (for some relevant references see Haw, 2009, section IIB, Kaye, 1982, section I)[6]. The defender of the risk minimisation picture could, of course, simply put such conventions down to prejudice as well – an unreflective, knee-jerk preference for testimonial over statistical evidence. Once again,

[5] We might think of this is an instance of what is sometimes called *statistical syllogism*: If I have evidence to the effect that a is a G and that the proportion of Gs that are F is x and have no further evidence relevant to whether a is an F, then the evidential probability, for me, that a is an F will be equal to x.

[6] This example has been the subject of a good deal of discussion within legal theory (for a review, see Ho, 2008, chapter 3, section 2.6), and appears to trace back to a genuine civil case – *Smith vs Rapid Transit Inc* 317 Mass. 469, 58 N.E. 2d 754 (1945). Many have concluded, on the basis of examples such as this, that standards of legal justification cannot be understood in purely probabilistic terms – a direct analogue, in effect, of the view that I shall defend about standards of epistemic justification.

this may turn out to be the right thing to say – but I shall experiment, in the next section, with a possible way of giving these intuitions and conventions more credit.

III. INTRODUCING NORMIC SUPPORT

Turning back to the laptop example, consider again the relationship between the proposition that (P_1) Bruce's laptop is displaying a blue background, the proposition that (P_2) my laptop is displaying a blue background, and my available evidence. Clearly, my evidence entails neither P_1 nor P_2. It would be perfectly possible for E_1 and E_2 to both be true while P_1 and P_2 are both false. Notice, though, that if E_1 is true and P_1 is false, then this would appear to be a circumstance crying out for *explanation* of some kind. If I visually perceive that Bruce's laptop is displaying a blue background when in fact it is not, then there has to be some explanation as to how such a state of affairs came about. Possible explanations have already been floated above – perhaps I'm hallucinating, or have been struck by colour blindness, or am subject to an optical illusion etc. It can't be that I *just misperceive* – there has to be more to the story.

The circumstance in which E_1 and P_1 are both true, we might say, is *explanatorily privileged* over the circumstance in which E_1 is true and P_1 is false. E_2 and P_2, however, do not appear to stand in this relationship. Although it would be very unlikely for P_2 to be false while E_2 is true, this is not something that would require special explanation. All of the random numbers that might be generated by my laptop are on an *explanatory par*. The occurrence of the one 'red' number does not require any more explanation than the occurrence of the 999 999 'blue' numbers. This, indeed, is part of what is involved in conceiving of a process as genuinely *random*.

If my belief that Bruce's laptop is displaying a blue background turned out to be false then, given the evidence upon which it is based, there would have to be some explanation for my error – either in terms of perceptual malfunction or disobliging features of the environment. If my belief that my laptop is displaying a blue background turned out to be false then, given the evidence upon which it is based, there need *not* be any available explanation for the error – the buck, as it were, may simply stop with me and the way that I chose to form my belief.

We can draw a similar contrast between the two kinds of evidence at play in the bus example. In scenario one, if it turned out that the bus involved was not owned by the Blue-Bus company, *in spite* of the eye-witness testimony, then there would have to be some accompanying explanation – the eye witness was lying, or was momentarily distracted, was mistaken about what Blue–Bus buses look like etc. It can't *just so happen* that the testimony was wrong. In scenario two, though, it could *just so happen* that the bus involved was not a Blue-Bus bus in spite of the fact that 95% of the buses in the area were. While this would, in a sense, be surprising, given the proportions involved, it clearly wouldn't demand any further special explanation.

The idea that *normalcy* is purely a matter of statistical frequency or propensity is, undeniably, an attractive one. Adopting it, though, forces us to give up on another attractive idea – namely, that normal conditions require *less explanation* than abnormal conditions do. Sometimes when we use the term 'normal' – when we say things like 'It's normal to be right handed' – we might be making a straightforward claim about statistical frequency. Other times – when we say things like 'Tim would normally be home by six' or 'When I turn my key in the ignition, the car normally starts' – part of what we are trying to express, I believe, is that there would have to be some satisfactory *explanation* if Tim wasn't home by six or the car wasn't starting.

In this sense of 'normal' it could be true that Tim is normally home by six, even if this occurrence is not particularly frequent. What *is* required is that exceptions to this generalisation are always explicable *as* exceptions by the citation of independent, interfering factors – his car broke down, he had a late meeting etc. If this condition is met, then the best way to explain Tim's arrival time each day is to assign his arrival by six a *privileged* or *default* status and to *contrast* other arrival times with this default (see Pietroski and Rey, 1995).

This may well be possible even if the number of occasions on which Tim arrived home by six was outweighed by the number of occasions on which he arrived home later. Suppose Tim is significantly delayed, day after day, by protracted roadworks on his usual route home and, were it not for the roadworks, he would always arrive home by six. There's a sense of 'normal' on which it remains true that Tim *normally* arrives home by six – we could imagine saying 'Tim would normally arrive home by six, but these blasted roadworks just keep delaying him day after day'.

We cannot explain everything at once – we need to abstract away from certain things in order to expose underlying patterns. This is how the idea of explanatory normalcy arises. The notion of explanatorily normal conditions can be usefully compared to the familiar idea of an *idealised* or simplified model of a potentially complex actual phenomenon. In any case, I won't attempt here to give a full, philosophically satisfactory account of explanatory or idealised normalcy. My immediate aim is more modest – to present, in rough outline, an alternative *kind* of normalcy that stands apart from the statistical conception that can so easily entrance us.

Given my evidence, P_2 would frequently be true, but P_1 would *normally* be true. That is, given my evidence, if P_1 turned out to be false, then this would demand explanation via some independent interfering factor. We might say, in this case, that my evidence *normically supports* P_1. Given my evidence, if P_2 turned out to be false, then no such explanation need be available. My evidence does not normically support P_2.

Once again, we can draw a similar contrast between to the two kinds of evidence mentioned in the bus case. Given that 95% of the buses operating in the area of an accident were Blue-Bus buses, it would frequently be true that the bus involved in the accident was a Blue-Bus bus. Given that an eye witness testified that the bus involved was a Blue-Bus bus, it would normally be true that the bus involved was a Blue-Bus bus. Our evidence in scenario one normically supports the conclusion that a Blue-Bus bus was involved, while our evidence in scenario two does not.

The distinction between these different sorts of evidential support might fruitfully be compared to the distinction between statistical generalisations and *ceteris paribus* generalisations widely accepted in the philosophy of science (see, for instance, Millikan, 1984, pp5, 33-34, Pietroski and Rey, 1995, 1.2). It might also be compared to the distinction, widely recognised in the philosophy of language, between generics that contain frequency adverbs – like ‘As are frequently B’, ‘As are typically B’ – and generics that are ‘unmarked’ – generics of the form ‘As are B’ (see, for instance, Leslie, 2008). These comparisons may be particularly apt if we are inclined to understand evidential probability along externalist lines. On the externalist conception of evidential probability, to say that a piece of evidence probabilifies a given proposition is, quite literally, to make a kind of statistical generalisation. To claim that a piece of evidence normically supports a proposition is at least *close* to making a kind of ceteris paribus generalisation.

I have characterised a support relation that evidence may bear to a proposition that does not simply reduce to probabilification. And it is not difficult to appreciate, at least in a rough and ready way, why this relation might have some connection with epistemic justification. If one believes that a proposition P is true, based upon evidence that normically supports it then, while one's belief is not assured to be true, this much is assured: If one's belief turns out to be false, then the error has to be *explicable* in terms of disobliging environmental conditions, deceit, cognitive or perceptual malfunction or some other interfering factor. In short, the error must be attributable to *mitigating circumstances* and thus *excusable*, after a fashion. Errors that do not fall into this category are naturally regarded as errors for which one must bear *full responsibility* – errors for which there is no excuse. And if error could not be excused, then belief cannot be permitted.

What I propose is that, in order for one to have justification for believing a proposition, it must be normically supported by one's evidence. When one classifies a belief as justified, one is effectively committed to the claim that, if the belief is *not* true, then this failure will be independently explicable in terms of some identifiable interfering factor. To borrow a turn of phrase used by Pietroski and Rey (1995, pp84), the notion of justification answers to the need to idealise in a complex world, not the need to describe a chancy one. The reason that I have justification for believing that Bruce's laptop is displaying a blue background, but lack justification for believing that my laptop is displaying a blue background is that the former proposition is normically supported by my evidence while the latter proposition is not. In an ideal world, given my evidence, Bruce's laptop would be displaying a blue background. But mine need not be.

The idea that there is a normic support condition on justification may seem reminiscent of a theory sometimes known as *abductivism* or *explanationism*. According to the explanationist, in order for one to have justification for believing a proposition, the truth of that proposition must, roughly speaking, be a part of the best or simplest explanation of one's relevant evidence (Conee and Feldman, 2008, section 3.5)[7]. What I am proposing here, though, is not a kind of explanationism. It is possible for a piece of evidence E to normically support a proposition P – for the circumstance in which E and P are both true to be explanatorily privileged over the circumstance in which E is true and P is false – without P providing any sort of explanation *for* E. Consider the following example (due to Byerly, forthcoming): Suppose it's a warm, calm day and I'm at the golf course facing a six foot putt to the hole. I strike the ball well – it appears to be rolling directly towards the hole with just the right amount of momentum to reach it. I believe that the ball will soon fall into the hole and it's plausible that my present evidence normically supports this. If the ball does not fall into the hole, in spite of my present evidence, then there will have to be some explanation as to why – a sudden gust of wind blows the ball off course, it strikes something on the green that diverts it or slows it down etc. But the proposition that the ball will soon fall into the hole doesn't *explain* my current evidence. That is, the proposition that the ball will fall into the hole at some future time doesn't explain my visual experiences *now*. It is controversial

[7] A number of epistemologists have defended what might be regarded as explanationist accounts of *knowledge* according to which, roughly speaking, in order for one to know that P, P must enter into the best explanation of one's belief that P (Jenkins, 2008, chap. 3, Goldman, 1988, chap. , Rieber, 1998). The concerns about explanationist accounts of justification that I outline in the body text may carry over to views of this kind, but I won't explore this here.

whether propositions about the future can ever enter into explanations of facts about the present or past. But there is no particular obstacle to propositions about the future being normically supported by facts about the present[8].

The condition that P be normically supported by one's evidence is not the same as the condition that P provide the best explanation for one's evidence. To infer, from a set of premises, a conclusion that they normically support is not to draw an *abductive* inference, as conventionally conceived – it is to draw a defeasible inference of a different kind. The claim that there is a normic support condition on justification is, of course, compatible with there being an explanationist condition as well. Examples such as the above, however, incline me to doubt that such a condition is, in general, necessary for justification – though it may be met in many cases. As I strike the ball and watch it roll in the right direction with just the right amount of momentum, surely my belief that it will soon fall into the hole is a justified one.

IV. NORMAL WORLDS

Suppose that possible worlds can be compared for their normalcy. Normic support could, then, be understood in terms of variably restricted quantification over worlds: A body of evidence E normically supports a proposition P just in case P is true in all the most normal worlds in which E is true. Alternately, E normically supports P just in case the most normal worlds in which E is true and P is true are more normal than any world in which E is true and P is false. This opens up a natural way of modelling normic support comparisons: A body of

[8] In order for this kind of example to work, it is not, in fact, necessary that the proposition in question concern the future – all that is really required is that my acquisition of the relevant evidence predate the time at which the truth value of the proposition in question is settled. If, after starting to watch the ball, I close my eyes and wait then, after a short while, I have evidence normically supporting the proposition that the ball is currently in the hole – namely, the perceptual experiences that are still fresh in my mind. But the proposition that the ball is currently in the hole doesn't explain this evidence.

evidence E normically supports a proposition P *more strongly* than it normically supports a proposition Q just in case the most normal worlds in which E is true and Q is false are more normal than any world in which E is true and P is false[9].

In modelling normic support in this way, I am effectively assuming that there will be maximally normal worlds in which P is true, for any contingent proposition P. That is, I'm assuming that we cannot have a situation in which, for every P-world, there is a more normal P-world. For simplicity, I shall leave this assumption in place. If we were to relax it though, we would need to alter the characterisation of normic support as follows: A body of evidence E normically supports a proposition P just in case some world at which E is true and P is true is more normal than any world at which E is true and P is false. A body of evidence E normically supports a proposition P more strongly than it normically supports a proposition Q just in case some world in which E is true and Q is false is more normal than any world in which E is true and P is false.

On my account, in order for one to have justification for believing a proposition P, it is necessary that one's body of evidence E normically support P – it is necessary that the most normal worlds in which E is true are worlds in which P is true. The probability of P given E can reach any level (short, perhaps, of 1) without this condition being met. Thus, the

[9] It is a consequence of this definition that E normically supports P iff E normically supports P more strongly than ~P. A corresponding principle for justification might be thought counterintuitive: One has justification for believing P iff one has more justification for believing P than ~P. It seems natural to think that, in order for one to have justification for believing P it is *not* enough that one simply have more justification for believing P than ~P. More generally, given this definition of comparative normic support, it will turn out that if E normically supports P more strongly than any other proposition Q, then E must normically support P. Once again, the corresponding principle for justification seems questionable: If one has more justification for believing P than Q then one has justification for believing P. Since normic support is proposed here merely as a *necessary* condition on justification, these latter principles won't be derivable from the former principles. I hope to discuss these issues elsewhere.

probability of P given E can reach any level (short, perhaps, of 1) without one having justification for believing P[10].

The following thought fits very naturally with this: In order for one to have *more* justification for believing a proposition P than a proposition Q, it is necessary that one's evidence E normically support P more strongly than Q – it is necessary that any world in which E is true and P is false be less normal than the most normal worlds in which E is true and Q is false. The probability of P given E can exceed the probability of Q given E without this condition being met. The probability of P given E can exceed the probability of Q given E without one having more justification for believing P than Q.

One thing that we might immediately notice about the notion of normic support, as modelled here, is that it will be closed under multiple premise deductive consequence. If E normically supports P and E normically supports Q and P and Q together entail R then it follows that E normically supports R. If E normically supports P, then the most normal E-worlds are P-worlds. If E normically supports Q then the most normal E-worlds are Q-worlds. If P and Q jointly entail R then all $P \wedge Q$-worlds are R-worlds. It follows straightforwardly that all the most normal E-worlds are R-worlds. The argument can be easily generalised for any number of premises.

[10] I am inclined to think that there will be possible cases in which one lacks justification for believing even propositions that have a probability of 1, given one's evidence. The fact that a proposition has an evidential probability of 0 need not imply that it is ruled out by, or impossible in light of, one's evidence. Suppose my evidence is enough to situate the speed of light within a particular interval, but leaves it equally likely that any of the values in this interval represents the true speed (our best evidence regarding the speed of light is, I take it, just like this). For any integer n, this intertval could be divided into n mutually exclusive and jointly exhaustive sub-intervals such that it is equally likely, given my evidence, that the true value lies within any one of these intervals as any other. The evidential probability that the speed of light is any particular value x within the interval must be less than 1/n for any integer n – that is, it must be equal to 0, unless we are willing to countenance infinitesimals. Although the probability that the speed of light is not equal to x is 1, this is perfectly consistent with there being possible worlds in which it is equal to x. As such, my evidence does not entail the proposition that the speed of light is not x and, provided some of these worlds are sufficiently normal, need not normically support it either.

Relatedly, given the possible worlds model that has been employed here, it can be shown that the conclusion of a multiple premise deductive inference cannot be less normically supported than each of its premises. Notice, first, that the normic support strength that a body of evidence E affords the conjunction $P \wedge Q$ must be equal to the normic support strength afforded to the least supported conjunct. The most normal worlds in which E is true and $P \wedge Q$ is false must either be the most normal worlds in which E is true and P is false or the most normal worlds in which E is true and Q is false – whichever are more normal. Notice, second, that if $P \wedge Q$ entails R, such that all $P \wedge Q$-worlds are R-worlds, then the normic support strength that E affords R can be no lower than the normic support strength that it affords $P \wedge Q$. Once again, the point clearly generalises beyond the two premise case.

Since normic support is proposed only as a necessary condition for justification, it won't follow from these observations that *justification* is closed under multiple premise deductive consequence. Whether we have this result will depend upon what else, if anything, we take justification to require. The laptop example purports to show that high evidential probability is not sufficient for justification but does not purport to show that high evidential probability is not necessary for justification. My arguments are quite compatible with the following 'hybrid' account of justification: One has justification for believing P only if P is normically supported by one's evidence and P is likely, given one's evidence, to be true. On this view, multiple premise closure will fail (though its failure will not be due to the normic support condition). Though it lacks a certain overall elegance I have not, as admitted, provided any reasons to think that this view is incorrect – and won't attempt to do so here.

V. OBJECTIONS AND REPLIES

As I suggested above, my description of the laptop example and, in particular, of the evidence available to Bruce and myself, is potentially contentious. According to the *knowledge account* of evidence (Williamson, 2000, chap. 9) one's evidence is equal to one's knowledge – that is, one's body of evidence consists of all and only the propositions that one knows. One who is impressed by this account might object to my description as follows: Once I *see* that Bruce's laptop is displaying a blue background, I come to *know* that it is and, according to the knowledge account of evidence, this suffices for the proposition to qualify as part of my body of evidence. My evidence, then, will include not just E_1 and E_2, but *P_1 as well*. In this case, the probability of P_1, given my evidence, will be 1 and, thus, *will* exceed the probability P_2, given my evidence, contrary to what I have claimed[11].

There is, I think, something dissatisfying about the way in which the knowledge account would have us treat this case – though it's difficult to put one's finger on precisely what it is. As I noted above, provided we fill in the details of the example in the most natural way, my belief that P_1 will qualify as a piece of knowledge. And part of what is puzzling about the example is *why this should be so* – why my belief that P_1 should qualify as knowledge, while my belief that P_2 seemingly never could. Pointing out that P_1 is more

[11] Epistemological disjunctivism about perceptual experience (McDowell, 1982, Pritchard, 2012) may make the very same prediction. According to the epistemological disjunctivist, roughly speaking, there is a fundamental epistemic asymmetry between veridical and non-veridical perceptions, such that the former put one in 'direct contact' with a fact in the world, while the latter do not. As such, it is a mistake to think that the evidence provided by veridical experience is no better than that available when one is subject to an illusion or hallucination. This view poses no particular threat to my construal of the laptop example, for much the same reason that the knowledge account of evidence poses no threat. As I shall show, it isn't essential to the example that my perceptual experience of Bruce's laptop even *be* veridical – we could perfectly well suppose it to be a hallucination. It might also be pointed out that examples of this sort need not involve direct perceptual evidence at all – as the bus case illustrates, we can, for instance, build an example around a contrast between statistical and testimonial evidence.

strongly supported by my evidence than P_2 looks like the beginnings of an answer to this question – but *not*, it seems, if the knowledge account of evidence is correct. According to the knowledge account, the only reason that P_1 is more strongly supported by my evidence than P_2 is *because* my belief that P_1 is taken to qualify as knowledge while my belief that P_2 is not. Any attempt to explain the knowledge disparity in terms of an underlying evidential disparity quickly re-implicates the very knowledge disparity to be explained.

It may be that these remarks just betray an unwillingness to enter into the spirit of 'knowledge-first' epistemology – but we needn't pursue this matter further here. In order to answer this objection, we need only point out that the force of the laptop example does not depend in any way upon the *actual background colour* – either of Bruce's laptop or of mine. Indeed, my original description of the case left it open what colour background the two laptops were actually displaying. Suppose that the background colour of both laptops is, in fact, red and I *really do* hallucinate when I look at Bruce's (unbeknownst to me of course). In this case, even the proponent of the knowledge account would, presumably, have to restrict my relevant evidence to something along the lines of E_1 (see Williamson, 2000, pp198). But the relevant intuitions are unchanged. Even though I'm now wrong about the background colour of both laptops, it still seems that I have *better* justification for believing that Bruce's background is blue than I do for believing that my background is blue. It would still be more legitimate for me to assert or act upon the former proposition than the latter. And the former belief is still a *candidate* for knowledge, while the latter is not.

The knowledge account can, in fact, be made to generate the same evidential predictions without our needing to suppose that I'm mistaken about the background colour of Bruce's laptop. We could imagine instead that I'm Gettiered when I set eyes on Bruce's laptop. Suppose there was some appreciable risk that I might have suffered a colour hallucination at that moment, even though this didn't eventuate and my perceptual experience was perfectly veridical. So long as the proponent of the knowledge account buys into the standard verdict about cases like this, he will, as before, have to restrict my evidence to something along the lines of E_1 as this is the only relevant knowledge in the vicinity. But whether one is Gettiered with respect to a proposition should have no effect on whether one has justification for believing it – or so the conventional thinking would have it, at any rate.

The knowledge account of evidence may not be the most natural fit with my original description of the evidence available to me in the laptop example – but, as this discussion shows, it is not inconsistent with this description. On the contrary, provided the details of the example are filled out in an appropriate way, the knowledge account can be made to explicitly sanction the claim that my relevant evidence consists of just E_1 and E_2.

While a proponent of knowledge account of evidence might worry that I'm underestimating the evidence available to me, one who is attracted to a narrowly psychologistic account evidence might worry, instead, that I'm *overestimating* this evidence. On a psychologistic account of evidence, roughly speaking, one's evidence is supposed to be limited to propositions about one's own mental states – propositions describing one's perceptual experiences, apparent memories, emotions etc. The suggestion that E_2 is literally a part of my evidence seems not to be consistent with this. E_2 after all, is a proposition about the world, and one that rests upon other propositions about the world – that my laptop is currently on and that the random number generator worked etc. If the psychologistic conception of evidence is correct then it couldn't literally be part of my evidence that my laptop is on or that the random number generator worked.

Let S comprise the various worldly conditions upon which E_2 rests. In order for the laptop example to engage the risk minimisation conception, it is not necessary that S literally be a *part* of my evidence or even that it be certain given my evidence. In order for the example to work, it is enough that the evidential probability of S, for me, be *greater* than the evidential probability of P_1. Let Pr be my evidential probability function once I lay eyes upon Bruce's laptop. By the theorem of total probability $Pr(P_2) = Pr(P_2 \mid S).Pr(S) + Pr(P_2 \mid {\sim}S).Pr({\sim}S)$ (provided $0 < Pr(S) < 1$). $Pr(P_2 \mid S)$ is, naturally, equal to 0.999999 – the evidential probability that my laptop is displaying a blue background *given that* it is on and the random number generator worked etc. will be equal to 0.999999. And we could, of course, push this value as close to 1 as we like, by suitably adjusting the details of the random number generator. Provided, then, that $Pr(S) > Pr(P_1)$, we can find a value for $Pr(P_2 \mid S)$ such that $Pr(P_2 \mid S). Pr(S) > Pr(P_1)$. It follows that $Pr(P_2) > Pr(P_1)$, as required.

If one is uncomfortable with the idea that S could be certain, given my evidence, then one could imagine instead that S is more likely, given my evidence, than P_1. But one who is attracted to a psychologistic conception of evidence might be resistant even to this suggestion. S, as noted above, comprises a number of worldly conditions – that my laptop is currently on and that the random number generator worked. Could my psychological evidence really make it more likely that these conditions are met than that Bruce's laptop is displaying a blue background? For what it's worth, I think the answer should be 'yes', provided the details of the example are filled out in the right way. It would be very odd for the psychologistic account of evidence (or any account of evidence) to place some a priori limit, strictly less than 1, on how evidentially likely a proposition like S could be. Be that as it may, though, I think it is possible to side-step this kind of position.

Let S, once again, comprise all of the worldly conditions upon which E_2 rests – that my laptop is currently on and that the random number generator worked etc. Whatever one takes the evidential probability of S or of P_2 to be, one should concede that the evidential probability of $S \supset P_2$ is *at least* 0.999999. That is, the evidential probability that my laptop is displaying a blue background *if* it is on and the random number generator worked, will be at least 0.999999. By the inclusion-exclusion rule, $Pr(S \supset P_2) = Pr(\sim S) + Pr(P_2) – Pr(\sim S \wedge \sim P)$. By the theorem of total probability, $Pr(S \supset P_2) = Pr(\sim S) + Pr(P_2 \mid S).Pr(S) + Pr(P_2 \mid \sim S).Pr(\sim S) – Pr(\sim S \wedge P_2)$ (provided $0 < Pr(S) < 1$). Even if we concede that $Pr(P_2 \mid \sim S) = 0$ – that there is *no chance* that my laptop is displaying a blue background if S is false – we still have it that $Pr(S \supset P_2) = Pr(\sim S) + Pr(P_2 \mid S).Pr(S)$. Since $Pr(S) = 1 – Pr(\sim S)$ and $Pr(P_2 \mid S) = 0.999999$, it follows that $Pr(S \supset P_2) = Pr(\sim S) + 0.999999.(1 – Pr(S)) = 0.999999 + 0.0000001.Pr(\sim S) \geq 0.999999$.

And yet, in the case described, a belief that $S \supset P_2$ would seem to be in *no better epistemic shape* than a belief that P_2. If I were to believe that my laptop is displaying a blue background, if it is on and the random number generator worked, then it would be very natural to describe this as a *presumption*. Such a belief, furthermore, would not constitute knowledge even if it turned out to be true. If someone were to ask me what colour background my laptop is displaying I shouldn't reply by baldly asserting 'If it's on and the random number generator worked then it's blue'. Rather, I ought to be more circumspect and utter a probabilistic conditional like 'If it's on and the random number generator worked then it's overwhelmingly likely to be blue'.

It's no more plausible that I have justification for believing $S \supset P_2$ than that I have justification for believing P_2. As such, given that the evidential probability of $S \supset P_2$, for me, exceeds the evidential probability of P_1, the example will continue to engage with the risk minimisation conception. We can simply put aside the question of how evidentially likely P_2 is and let the example be driven by the contrast between P_1 and $S \supset P_2$.

A final objection to my description of the laptop example takes issue not with my description of the available evidence but, rather, with my assumptions as to how this evidence serves to confer evidential probability values upon P_1 and P_2. The idea that propositions should always have precise evidential probability values is one that is often challenged – particularly when the evidence in question is not of an overtly statistical kind (see, for instance, Levi, 1985, Sturgeon, 2008). While my evidence in favour of P_2 is of a broadly statistical kind, my evidence in favour of P_1, of course, is not. While we have no particular trouble accepting that the probability of P_2 given my evidence is equal to 0.999999, the selection of any precise value for the probability of P_1 given my evidence is likely to strike one as artificial and contrived – there are many values that seem clearly wrong, but none that seem clearly right. And yet, haven't I effectively assumed, in my description of the laptop case, that P_1 does have some precise value, given my evidence – indeed a value less than 0.999999?

Just because we cannot find some natural or obvious precise value to assign to the probability of P_1, given my evidence, we should, I think, be cautious about concluding that it *has* no precise value. This inference seems to rest on the assumption that the probabilistic bearing of a body of evidence should always be transparent or accessible to us – but it isn't clear to me that such an assumption is viable in all cases. Certainly if we adopt an externalist construal of evidential probability the assumption would in general be false. On an externalist construal of evidential probability, the evidential probability of P_1, given my evidence, will be determined, in effect, by the frequency with which experiences of the sort described by E_1 are veridical in actual and relevant hypothetical scenarios. If an externalist is willing to accept that this frequency has a precise value, he should also be willing to accept that the evidential probability of P_1, given my evidence, has a precise value – albeit not one that we are in any immediate position to ascertain.

In any case, in order for the laptop example to have its desired effect, the evidential probability of P_1 need not take a precise value. All that is needed is the truth of a certain *comparative* claim – and, while it may seem unnatural for P_1 to be assigned a precise evidential probability value, it does not seem at all unnatural for P_1 to feature in certain evidential probability comparisons. Consider the following: 'It's more likely, given my evidence, that Bruce's laptop is displaying a blue background than a green one' and 'While it's likely for me that Bruce's laptop is displaying a blue background, it is not as likely as 99.9999%'. Both of these comparative claims seem clearly true, irrespective of whether we're inclined to think that P_1 has some precise evidential probability.

The final comparison is, of course, exactly what is needed in order for the laptop example to engage with the risk minimisation conception – assuming that 99.9999% is precisely how likely P_2 is, for me. In fact, if we wished, we could even do away with the idea that the evidential probability of P_2 takes a precise value. Pared down to its bare essentials, all that the example requires is the truth of the following comparison: P_2 is more likely, given my evidence, than P_1.

There are a number of ways of formally modelling imprecise probabilities – but the model that is, perhaps, most familiar to philosophers exploits a set of probability functions Γ – sometimes termed a *representor* – to capture the probabilistic bearing of a body of evidence (see, for instance, Levi, 1974, 1985, Van Fraassen, 1990, Hájek, 2003, Joyce, 2005, Weatherson, 2007). A proposition P will only have a precise evidential probability value in case all of the functions in one's representor agree on how likely the proposition is – if they all assign it the same value. Otherwise, P will be associated with a range of values. The composition of a representor is standardly constrained so as to ensure that these values form an interval[12].

The right way to apply such a model to the laptop case is, presumably, to have all of the functions in my representor assign a value of 0.999999 to P_2, but to have them assign a spread of values to P_1. So long as no function in my representor assigns P_1 a value as high as 0.999999, however, we will have the required comparison – namely, that the evidential probability of P_1 is less than 0.999999 and, indeed, less than the evidential probability of P_2. We could even, if we wish, accommodate the idea that some function in my representor assigns P_1 a value higher than 0.999999, so long as no function assigns it a value of 1. After all, we can simply alter the details of the random number generator so as to push the precise evidential probability of P_2 as close to 1 as we like.

[12] A representor is standardly required to be *convex* – that is, closed under the operation of taking weighted averages of probability functions. More formally, a representor Γ is convex just in case for every $Pr_x \in \Gamma$ and $Pr_y \in \Gamma$, $(\alpha Pr_x + (1 - \alpha)Pr_y) \in \Gamma$, for all α, $0 \leq \alpha \leq 1$. The weighted average of two probability functions is itself defined by taking the weighted average of the two values assigned to each proposition. More formally, for every proposition P over which Pr_x and Pr_y are defined $(\alpha Pr_x + (1 - \alpha)Pr_y)(P) = \alpha Pr_x(P) + (1 - \alpha)Pr_y(P)$. That the new function, so defined, is indeed a probability function can be easily proved. As can be clearly seen, the values assigned to a proposition by the functions in a convex representor must themselves be closed under the taking of weighted averages and, thus, form a real interval.

If some function in my representor *did* assign P_1 the value 1, then there would indeed be no way to make the crucial comparison come out as true. But, given the way in which evidential probability comparisons are evaluated on the present model, such a modelling choice would have exceedingly implausible consequences. A proposition P counts as less likely than a proposition Q, relative to representor Γ, just in case every function in Γ assigns P a lower value than Q – for all $Pr_x \in \Gamma$, $Pr_x(P) < Pr_x(Q)$ (see Joyce, 2005, section 2, Weatherson, 2007, section III). If we include in my representor a function that assigns P_1 the value 1 then, given my evidence, P_1 won't count as less likely than *any* other proposition and, accordingly, no other proposition will count as less likely than $\sim P_1$. But this doesn't seem to be an accurate reflection of my epistemic position. It may be unlikely, for me, that Bruce's laptop is not displaying a blue background, but it's surely even less likely, for instance, that I'll win 10 consecutive national lotteries or that I'm about to spontaneously combust or that Bruce's laptop is not displaying a blue background and there's an even number of blades of grass on my lawn.

In this paper I've argued that the risk minimisation conception of justification is flawed and have begun to sketch an alternative picture. In section I, I quoted some advice from Descartes – namely, that one should never believe that which is merely probable. In one sense, I concur with this. If the *only* thing that can be said in favour of a proposition is that it is probable, given one's evidence then, in my view, one would not be justified in believing it. In another sense, I, like any fallibilist, will reject Descartes' advice – in my view one can be perfectly justified in believing things that are less than certain. The compatibility of these two views owes to the fact that there are a range of ways in which evidence can make a proposition more than probable, though less than certain. I have outlined one such way here.

REFERENCES

Alston, W. (1988) 'An internalist externalism' *Synthese* v74, pp265-283

Bird, A. (2007) 'Justified judging' *Philosophy and Phenomenological Research* v74(1), pp81-110

BonJour, L. (2010) 'The myth of knowledge' *Philosophical Perspectives* v24(1), pp57-83

Byerly, T. (forthcoming) 'Explanationism and justified beliefs about the future' *Erkenntnis*

Carnap, R. (1950) *The Logical Foundations of Probability* (Chicago: University of Chicago Press)

Chisholm, R. (1957) *Perceiving* (Ithaca, NY: Cornell University Press)

Cohen, J. (1977) *The Probable and the Provable* (Aldershot: Gregg Revivals)

Cohen, S. (1988) 'How to be a fallibilist' *Philosophical Perspectives* v2, pp91-123

Comesaña, J. (2006) 'A well-founded solution to the generality problem' *Philosophical Studies* v129, pp27-47

Conee, E. and Feldman, R. (2008) 'Evidence' in Smith, Q. ed. *Epistemology: New Essays* (Oxford: Oxford University Press)

DeRose, K. (1996) 'Knowledge, assertion and lotteries' *Australasian Journal of Philosophy* v74, pp568-580

Descartes, R. [1628] *Rules for the Direction of the Mind* (Cambridge: Cambridge University Press, 1911)

Dudman, V. (1992) 'Probability and assertion' *Analysis* v52, pp204-211

Enoch, D., Fisher, T. and Spectre, L. (forthcoming) 'Statistical evidence, sensitivity and the legal value of knowledge' *Philosophy and Public Affairs*

Foley, R. (1987) *The Theory of Epistemic Rationality* (Cambridge, MA: Harvard University Press)

Foley, R. (1992) 'The epistemology of belief and the epistemology of degrees of belief' *American Philosophical Quarterly* v29, pp111-121

Foley, R. (1993) *Working Without a Net* (Oxford: Oxford University Press)

Foley, R. (2009) 'Beliefs, degrees of belief and the Lockean thesis' in Huber, F. & Christoph Schmidt-Petri, C. eds. *Degrees of Belief* (Springer)

Fumerton, R. (1995) *Metaepistemology and Scepticism* (Lanham, MD: Rowman and Littlefield)

Fumerton, R. (2005) 'The challenge of refuting scepticism' in Steup, M. and Sosa, E. eds. *Contemporary Debates in Epistemology* (Malden, MA: Blackwell Publishing)

Fumerton, R. (2004) 'Epistemic probability' *Philosophical Issues* v14(1), pp149-164

Goldman, Alan. (1988) *Empirical Knowledge* (Berkeley, CA: University of California Press)

Goldman, Alvin. (1967) 'A causal theory of knowing' *Journal of Philosophy* v64, pp357-372

Goldman, Alvin. (1986) *Epistemology and Cognition* (Cambridge, MA: Harvard University Press)

Hájek, A. (2003) 'What conditional probability could not be' *Synthese* v137(3), pp273-323

Hájek, A. and Smithson, M. (forthcoming) 'Rationality and indeterminate probabilities' *Synthese*

Harman, G. (1986) *Change in View* (Cambridge, MA: MIT Press)

Haw, R. (2009) 'Prediction markets and law: A sceptical account' *Harvard Law Review* v122, pp1217-1229

Hawthorne, J. (2003) *Knowledge and Lotteries* (Oxford: Oxford University Press)

Ho, H. (2008) *A Philosophy of Evidence Law: Justice in the Search for Truth* (Oxford: Oxford University Press)

Jeffrey, R. (1968) 'Probable knowledge' in Lakatos, I. ed. *The Problem of Inductive Logic* (Amsterdam: North Holland)

Jeffrey, R. (1970) 'Dracula meets wolfman: Acceptance vs. partial belief', in Swain, M. ed., *Induction, Acceptance, and Rational Belief* (Dordrecht: Reidel)

Jenkins, C. (2008) *Grounding Concepts* (Oxford: Oxford University Press)

Joyce, J. (2005) 'How probabilities reflect evidence' *Philosophical Perspectives* v19 pp153-178

Kaplan, M. (1981) 'A Bayesian theory of rational acceptance' *Journal of Philosophy* v78, pp305-330

Kaye, D. (1982) 'The limits of the preponderance of the evidence standard: Justifiably naked statistical evidence and multiple causation' *American Bar Foundation Research Journal* v7(2), pp487-516

Keynes, J. [1973] *A Treatise on Probability* (London: Macmillan press, 1921)

Kvanvig, J. (2007) 'Assertion, knowledge and lotteries' in Pritchard, D. and Greenough, P. eds. *Williamson on Knowledge* (Oxford: Oxford University Press)

Kyburg, H. (1961) *Probability and the Logic of Rational Belief* (Middleton: Wesleyan University Press)

Lempert, R. (2001) 'The economic analysis of evidence law: Common sense on stilts' *Virginia Law Review* v87

Leplin, J. (2009) *A Theory of Epistemic Justification* (Springer)

Leslie, S. (2008) 'Generics: Cognition and acquisition' *Philosophical Review* v117(1), pp1-47

Levi, I. (1974) 'On indeterminate probabilities' *Journal of Philosophy* v71, pp391-418

Levi, I. (1985) 'Imprecision and indeterminacy in probability judgments' *Philosophy of Science* v52(3), pp390-409

Levi, I. (2000) 'Imprecise and indeterminate probabilities' *Risk and Decision* v5, pp111-122

Lewis, D. (1980) 'A subjectivist's guide to objective chance' in *Studies in Inductive Logic and Probability* vol. II (University of California Press)

Lewis, D. (1996) 'Elusive knowledge' *Australasian Journal of Philosophy* v74(4), pp549-567

Maher, P. (1993) *Betting on Theories* (Cambridge: Cambridge University Press)

Makinson, D. (1965) 'The paradox of the preface', *Analysis* v25, pp205-7

McDowell, J. (1982) 'Criteria, defeasibility and knowledge' *Proceedings of the British Academy* v68, pp455-479

McGlynn, A. (forthcoming) 'Justification as would-be knowledge' *Episteme*

Millikan, R. (1984) *Language, Thought and Other Biological Categories* (Cambridge, MA: MIT Press)

Nelkin, D. (2000) 'The lottery paradox, knowledge and rationality' *Philosophical Review* v109, pp373-409

Pietroski, P. and Rey, G. (1995) 'When other things aren't equal: Saving "ceteris paribus" laws from vacuity' *British Journal for the Philosophy of Science* v46, pp81-110

Pollock, J. (1986) *Contemporary Theories of Knowledge* (Rowman and Littlefield)

Pollock, J. (1990) *Nomic Probability and the Foundations of Induction* (Oxford: Oxford University Press)

Pritchard, D. (2007) 'Sensitivity, safety and anti-luck epistemology' in Greco, J. ed. *The Oxford Handbook of Scepticism* (Oxford: Oxford University Press)

Pritchard, D. (2012) *Epistemological Disjunctivism* (Oxford: Oxford University Press)

Pryor, J. (2004) 'What's wrong with Moore's argument?' *Philosophical Issues 14: Epistemology*

Pryor, J. (2005) 'Is there immediate justification?' in Sosa, E. and Steup, M. eds. *Contemporary Debates in Epistemology* (Oxford: Blackwell)

Redmayne, M. (2008) 'Exploring the proof paradoxes' *Legal Theory* v14, pp281-309

Rieber, S. (1998) 'Scepticism and contrastive explanation' *Noûs* v32, pp189-204

Russell, B. (1948) *Human Knowledge: Its Scope and Limits* (London: Unwin)

Ryan, S. (1991) 'The preface paradox' *Philosophical Studies* v64, pp293-307

Ryan, S. (1996) 'The epistemic virtues of consistency' *Synthese* v109, pp121-141

Savage, L. (1972) *Foundations of Statistics* (New York: Dover)

Schiffer, S. (2003) *The Things We Mean* (Oxford: Oxford University Press)

Schurtz, G. (2001) 'What is "normal"? An evolution-theoretic foundation for normic laws and their relation to statistical normality' *Philosophy of Science* v68, pp476- 497

Smith, M. (2007) 'Ceteris paribus conditionals and comparative normalcy', *Journal of Philosophical Logic*, v36(1), pp97-121

Smith, M. (2009) 'Transmission failure explained' *Philosophy and Phenomenological Research* v79(1), pp164-189

Smith, M. (2010) 'What else justification could be' *Noûs* v44(1), pp10-31

Thompson, J. (1986) 'Liability and individualized evidence' in Parent, W. ed. *Rights, Restitution and Risk* (Cambridge, MA: Harvard University Press)

van Fraassen, B. (1990) 'Figures in a probability landscape' in Dunn, J. and Gupta, A. eds. *Truth or Consequences* (Amsterdam: Kluwer)

Vogel, J. (1990) 'Are there counterexamples to the closure principle?' in M. Roth and G. Ross (eds.), *Doubting: Contemporary Perspectives on Skepticism* (Dordrecht: Kluwer)

Vogel, J. (1999) 'The new relevant alternatives theory', *Philosophical Perspectives*, v13, pp155–80

Weatherson, B. (2005) ‘Can we do without pragmatic encroachment?’ *Philosophical Perspectives* v19, pp417-443

Weatherson, B. (2007) ‘The Bayesian and the dogmatist’ *Proceedings of the Aristotelian Society* v107, pp169-185

White, R. (2005) ‘Epistemic permissiveness’ *Philosophical Perspectives* v19(1), pp445-459

Williamson, T. (2000) *Knowledge and its Limits* (Oxford: Oxford University Press)