THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Between Probability and Certainty

Citation for published version:

Smith, M 2016, *Between Probability and Certainty: What Justifies Belief*. vol. 77, Oxford University Press, Oxford. <https://doi.org/10.1093/acprof:oso/9780198755333.001.0001>

Digital Object Identifier (DOI):

[10.1093/acprof:oso/9780198755333.001.0001](https://doi.org/10.1093/acprof:oso/9780198755333.001.0001)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

What Justifies Belief

2.1. 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 function¹. This is, perhaps, more

¹ 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 (Kolmogorov, 1933):

- (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 is taken to be infinite, and closed under countable disjunction, then (iii) might be strengthened to the following:

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

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

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 lies close to 1 and may be variable and/or vague² – but the details of this will matter little for the moment.

Some epistemologists have dissented from this picture and defended alternative views – the theories proposed by Sutton (2007), Bird (2007), Reynolds (2013) and Ichikawa (2014) have already been mentioned³. L.J. Cohen (1977, chap. 22, 1979), Dana Nelkin (2000) and Jarrett Leplin (2009) might also be added to this list. But, as noted in the last chapter, the majority of epistemologists do appear to accept that something along the lines of the risk minimisation conception is correct. And there is, undeniably, something very natural about it. After all, if justification does not require evidential certainty then what could it possibly require if not evidential likelihood? If justification does not demand the complete elimination of error risk then what could it possibly demand if not its minimisation? One could, when thinking in this way, almost get the impression that (a) and (b) are not substantial claims at all but, rather, serve to *define* a basic fallibilist notion of justification – a notion that comes before any substantial epistemic theorising.

It isn’t an essential part of the risk minimisation conception that evidential probabilities be viewed as classical – though this would be the orthodox stance. There are many alternative approaches; views that allow for infinitesimal probabilities or for interval-valued probabilities, axiomatisations that treat conditional probabilities as primitive, axiomatisations that embed nonclassical logics etc. The risk minimisation conception, as I understand it, is quite compatible with approaches such as these.

² Any evidential probability value that one chooses is bound to seem rather arbitrary as a threshold for justification. Making the threshold variable and/or vague is sometimes thought to help with this – though it’s not entirely clear how this is supposed to work. On the contrary, introducing a number of seemingly arbitrary values that the threshold can assume might be thought to compound the problem. Whatever the truth of the matter, I won’t make much of this ‘threshold problem’ here – but I will say a little more about it in future chapters.

³ Even Sutton, who holds that justification, strictly speaking, is identical with knowledge thinks that something like the risk minimisation conception often captures what we intend to *convey* with justification attributions. According to Sutton, when we say that one justifiably believes P often we are just speaking loosely and mean only that one justifiably believes that P is *probable* (Sutton, 2007, section 2.3), which may be true if the evidential probability of P , for one, is sufficiently high. I am inclined to think that these two sorts of judgment are not nearly so close as Sutton supposes.

The reflections of chapter 1 show us, at the very least, that (a) and (b) cannot really be trivialities. Whatever we make of the claim that knowledge and justification are normatively coincident goals, it is not something that is *trivially* false. The idea that beliefs can be appraised as permissible or impermissible, in a distinctively epistemic way, 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 captured by a platitude like this: One has justification for believing a proposition P iff one is permitted, epistemically, to believe P. Even if it is true that the minimisation of error risk is how we earn epistemic permission to believe things, it is not true by stipulation.

I’m inclined to think, however, that this is not true at all. In this chapter I shall outline an argument against the risk minimisation conception. I have already offered an argument of a kind against this conception – namely, one that proceeds from the assumption that knowledge and justification are normatively coincident goals. The argument I offer here will not use this as a premise – though it will appeal to it as a kind of supplementary consideration. As well as offering an argument against the risk minimisation conception, I shall begin the task of sketching a possible alternative picture. First, though, a little more stage setting is required.

I assume here that evidence is propositional – that one’s body of evidence consists of a stock of propositions or a conjunction of propositions (Williamson, 2000, section 9.5, Dougherty, 2011). 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 Williamson’s knowledge account of evidence, according to which one’s evidence consists of the propositions that one knows (Williamson, 2000, chap. 9), and with

more restrictive accounts that limit one's evidence to a proper subset of one's knowledge, such as the knowledge of one's own experiences and mental states (Lewis, 1996, Conee and Feldman, 2008, Swinburne, 2011)⁴.

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 – in particular, facts about *propensities* or *frequencies* (see Russell, 1948, chap. VI, Alston, 1988, section I, Plantinga, 1993, chap. 9, section I). On one sort of externalist view, the probability of a proposition P, given evidence E, will be determined by the frequency with which the kind of circumstance described by E is accompanied by the kind of circumstance described by P across actual and similar counterfactual circumstances. *Internalists*, on the other hand, conceive of evidential probability values as reflecting necessary, internal connections between evidence and hypotheses (see Keynes, 1921, Carnap, 1950, Kyburg, 1971, Moser, 1988, Fumerton, 1995, 2004, Williamson, 2000, chap. 10, Conee and Feldman, 2008, section 1.5). Once again, I take no stand on this issue here – but I will have more to say about it in future chapters.

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

⁴ Some epistemologists hold that one's evidence fundamentally consists of one's experiences and mental states *themselves*, rather than propositions about them. I have no particular objection to this way of speaking – but what I do think is that many of the relations that evidence bears to propositions have to be analysed, first and foremost, as inter-propositional relations. A proposition cannot literally be entailed or made probable, for example, by an experience *itself*, though it can be entailed or made probable by the *occurrence* of an experience – that is, by the proposition that the experience occurred (for discussion of this see Williamson, 2000, section 9.5, Dougherty, 2011). Propositions about one's experiences and mental states must be regarded as evidence in at least a derivative sense, even if we want to hold on to a stricter sense of the term reserved for the experiences and mental states themselves.

truth. I can regard it as *extremely* likely that, say, the number of stars in the universe is a composite number or that ticket #542 has lost the lottery without actually believing either of these things⁵. 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* (Kaplan, 1995, pp121). (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.

2.2. Problems for Risk Minimisation

One possible argument against the risk minimisation conception proceeds from the principle of *multiple premise closure*, mentioned in the last chapter. It's very plausible that deduction, of all things, should be an epistemically permissible way of expanding one's set of beliefs – that deduction, of all things, won't take me from beliefs that I am epistemically permitted to hold to beliefs that are epistemically impermissible. One way of attempting to make this rough intuition precise is the following: If one has justification for believing each of a set of

⁵ 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, 1993, chap. 4, 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. I shall have more to say about degrees of belief in section 8.4.

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 error 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 conclusion of a multiple premise deduction can inherit the error risk associated with each of the premises and, thus, have an evidential probability that dips below the threshold for justification, even if the evidential probability of each premise exceeds it. The risk minimisation theory would appear to be compatible with the weaker principle of *single premise closure*: If one has justification for believing a proposition, then one has justification for believing any proposition that deductively follows from it. If a deductive inference has but a single premise then there is no error risk to aggregate⁶.

The fact that the risk minimisation conception predicts the failure of multiple premise closure does, I think, constitute one kind of objection to it. But the dialectical situation concerning multiple premise closure is complex, and I won't try to press this objection yet – though I will return to it in chapter 4. The argument I will offer here involves a rather different

⁶ The closure principle described here concerns *propositional* justification. We could also formulate a related closure principle for doxastic justification: If one justifiably believes each of a set of propositions, and deduces a further proposition from them, then this further proposition is justifiably believed. To avoid certain counterexamples, we should specify that the deduction is careful and competent and that the relevant beliefs remain justified while performing it etc. Even when suitably qualified, we would expect this principle to fail on the risk minimisation conception for the reasons given in the body text, though a single premise version might, once again, remain secure.

The connection between the propositional and doxastic closure principles is very close – though mediated by certain further assumptions. If one justifiably believes each of a set of propositions then one must have justification for believing each of these propositions and believe each in a way that is based upon its justification. If these propositions deductively entail a further proposition then the propositional closure principle guarantees that one has justification for believing it. All that is then needed for the doxastic principle to be assured is the assumption that deduction preserves proper basing – if one (competently) deduces a conclusion from premises that are properly based on their justifications (throughout the deduction), then that conclusion will be properly based upon its justification.

prediction that the risk minimisation conception makes – a prediction concerning the force of evidence that is purely statistical in nature.

Consider the following 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 set eyes upon it, it simply *appears* to me to be displaying a blue background and I immediately 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):

(E₁) It visually appears to me that Bruce's laptop is displaying a blue background.

(E₂) 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* (perhaps 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, for instance, 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. To believe that Bruce’s laptop is displaying a blue background would be to believe in a way that is straightforwardly directed at knowledge. To believe that my laptop is displaying a blue background would be to believe in way that seems indifferent to knowledge.

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 scenarios – but presumably the likelihood, given evidence E_1 , that Bruce’s laptop is displaying a blue background would be nowhere 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_2^7 . 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. If this set-up were replicated again and again then, in the long run, we would expect my belief about Bruce's laptop to be false more often than my belief about my laptop.

The judgment that I lack justification for believing that my laptop is displaying a blue background is in considerable tension with the risk minimisation conception. This proposition is clearly *very likely*, given my evidence – by believing this proposition I would be running only a minute risk of error. If the risk minimisation theory is correct, then I could scarcely do better, epistemically, than to believe a proposition like this.

Provided the likelihood threshold t is set below 0.999999, the judgment that I lack justification for believing that my laptop is displaying a blue background will conflict with thesis (a) of the risk minimisation conception. 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. Try as we might, we will never find a likelihood threshold such that the proposition that Bruce's

⁷ 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 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.

background is blue lies on the right side of it and the proposition that my background is blue lies on the wrong side.

To refute thesis (b) of the risk minimisation, we needn't even make any categorical judgments about what I have justification for believing – it's enough that we make the *comparative* judgment that I have more or better justification for believing that Bruce's background is blue than I do for believing that my background is blue. Strictly, this is all that is required to refute the letter of the risk minimisation conception.

The laptop example prompts a cluster of judgments 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 judgments as confused or naïve. Perhaps we are simply accustomed to relying upon perception when it comes to such matters and suspending any scruples about its fallibility. Once we do reflect on the various ways that perception can go wrong, so this thought goes, these troublesome intuitions are exposed as a kind of prejudice, and my justificatory standing with respect to the two propositions in question no longer seems so different (see Steglich-Petersen, 2013, section 3.3). I'm not entirely convinced that this is the wrong thing to say – but I do suspect that it is.

Consider another example – one that is well known in legal theory and the philosophy of law, though less so in epistemology: Suppose 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 to the incident 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 kind⁸. The eye witness in scenario one could have suffered a hallucination or she could have fabricated her memory of the incident or she could have deliberately concocted a lie in order to smear the Blue-Bus company etc. These possibilities may not be likely, but if we were forced to come up with some estimate of the probability that the bus involved really was a Blue-Bus bus, given the witness testimony, it's doubtful that we would go quite as high as 95% – this would seem overly trusting. But this, of course, is precisely how probable the proposition is, given the strong statistical evidence available in scenario two⁹.

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 caused the incident. Indeed, if this is our only relevant evidence, it would seem reasonable for us to assert and to act upon this conclusion. But what about scenario two? Should I really believe that the bus involved was a Blue-Bus bus on the grounds that 95% of the buses in the area on the day in question were Blue-Bus buses? Should I go around announcing that the bus involved was a Blue-Bus bus? Should I take steps against the company – boycott their buses, picket their offices etc? If my *only evidence* is that 95% of the buses in the area on the day in question were Blue-Bus buses then to take such steps would surely be unjust (see Redmayne, 2008, Enoch, Fisher and Spectre, 2012, Buchak, 2014). On the risk minimisation conception, though, any unwillingness to believe or to act in this scenario is very puzzling. After all, the proposition

⁸ See, for instance, Loftus (1996).

⁹ If one thought that 95% really is a reasonable estimate, in scenario one, of how likely it is that the bus involved was a Blue-Bus bus, we could simply increase the proportion of Blue-Bus buses in scenario two to restore the intended overall structure of the example. The probability value in scenario two is derived using 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.

that the bus involved was a Blue-Bus bus is made *very likely* by the statistical evidence I possess. By believing this proposition I would only be running a very small risk of error – I would have managed the risk of error in an almost exemplary way.

Questions about the force of purely statistical evidence are not, of course, purely academic – and it is not just theorists who have converged on these answers. Under prevailing legal practice, across a range of jurisdictions, eye witness testimony of the sort available in scenario one could serve as a legitimate basis for a finding to the effect that the bus involved was a Blue-Bus bus – and for a finding of liability in a civil trial. In contrast, it would not be considered legitimate to base a finding of liability upon the kind of statistical evidence available in scenario two. Indeed, the general reluctance of courts, and individuals, to base verdicts of guilt or liability on evidence that is purely statistical in nature is well established¹⁰. The defender of the risk minimisation picture could, of course, simply put this pattern down to prejudice as well – some kind of unreflective, knee-jerk preference for testimonial over statistical evidence. Once again, 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.

¹⁰ The Blue-Bus example in fact traces back to a genuine civil case – *Smith vs Rapid Transit Inc* 317 Mass. 469, 58 N.E. 2d 754 (1945). The trial court, in this case, found that the imposition of liability could not rest upon the kind of purely statistical evidence supplied by the complainant – a verdict upheld, on appeal, by the Supreme Judicial Court of Massachusetts. These verdicts are very much in keeping with a broader pattern in which purely statistical evidence has been judged insufficient to meet legal burdens of proof. Other cases that are sometimes cited in this regard include *Virginia & S.W. Ry. Co. v. Hawk* 160 F 348, 352 (6th cir, 1908), *Evans vs Ely* 13 F.2d 62, 64 (3rd Cir, 1926), *Commercial Standards Insurance Co. vs. Gordon Transports Inc.* 154 F.2d 390, 396 (1946), *People vs. Collins* 438 P.2d 33 40-41 (Cal. 1968), *State vs. Carlson* 267 N.W. 170, 179 (Minn. 1978), *United States vs. Shonubi* 103 F 3.d 1085 (2d Cir, 1997), *R vs. Watters* Court of Appeal, Criminal Division (19th October 2000). For further relevant references, research and discussion see Kaye, (1982, section I), Wells (1992) and Haw (2009, section IIB). Many have concluded, on the basis of the case law, 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.

2.3. Introducing Normic Support

Turning back to the laptop example, consider again the relationship between the proposition (P_1) that Bruce's laptop is displaying a blue background, the proposition (P_2) that 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 the 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 Blue-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 hallucinating or lying or had a fabricated memory. Whatever the truth of the matter, it can't *just so happen* that the testimony was wrong – there has to be more to the story. 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, he had to detour around roadworks 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 – first by car trouble, then by a late meeting, then by roadworks etc. – but, were it not for these interfering factors, 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 be here by six, but he's just had a bad run lately!'

Say that a body of evidence E *normically supports* a proposition P just in case the circumstance in which E is true and P is false requires more explanation than the circumstance in which E and P are both true (Smith, 2010a). Given my evidence in the laptop example, the circumstance in which Bruce's laptop is not displaying a blue background would require more explanation than the circumstance in which it is. In contrast, the situation in which my laptop is not displaying a blue background, as unlikely as that might be, would not require more explanation than the circumstance in which it is. My evidence normically supports the proposition that Bruce's laptop is displaying a blue background, but does not normically support the proposition that my laptop is displaying a blue background.

Turning to the Blue-Bus example, if we have evidence to the effect that an eye witness testified that the bus involved in the incident was a Blue-Bus bus, then the circumstance in which it wasn't a Blue-Bus bus requires more explanation than the circumstance in which it was. The testimonial evidence normically supports the proposition that the bus involved was a Blue-Bus bus. If, on the other hand, we have evidence to the effect that 95% of the buses

operating in the area are Blue-Bus buses, then the circumstance in which the bus involved wasn't a Blue-Bus bus requires no more explanation than the circumstance in which it was. Given that 95% of the buses operating in the area were Blue-Bus buses, it would *frequently* be true that the bus involved in the incident 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.

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. As discussed above, 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. I'll have more to say about such comparisons in chapter 6.

I have characterised an evidential support relation that demands more than probability, but less than certainty. 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 etc. In short, the error must be attributable to *mitigating circumstances* of some kind 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 committed to the claim that, if the belief fails to be 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.

2.4. Normal Worlds

Suppose that possible worlds can be compared for their normalcy¹¹. Normic support could, then, be modelled 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, we might say that E normically supports P just in case there is a world at which E is true and P is true which is more normal than any world at which E is true and P is false. In fact, these two characterisations are not quite equivalent – the former assumes

¹¹ The idea that possible worlds or states of affairs might be ordered with respect to their normalcy has been proposed before. It has been explored in connection with conditional logics for defeasible reasoning (Delgrande, 1987, Boutilier, 1994, Boutilier and Becher, 1995), in connection with counterfactual conditionals (Gundersen, 2004) and in connection with conditional analyses of causation (Menzies, 2004). I shall return to some of these applications of the idea in chapter 6.

that there will be maximally normal worlds in which E is true, for any conceivable body of evidence E. I'll make this assumption for now – but will return to it, particularly in section 8.1. We also have a natural way of modelling normic support comparisons: A body of evidence E normically supports a proposition P *more strongly* than it normically supports a proposition Q just in case there is a world at which E is true and Q is false which is more normal than any world at which E is true and P is false¹².

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 all the most normal worlds at which E is true are worlds at which P is true. The probability of P given E can reach any level, short perhaps of 1, without this condition being met. If the probability of P given E is less than 1, there will be possible worlds at which E holds and P does not. If these are amongst the most normal worlds at which E holds, then E will not normically support P. On my account, the probability of P given E can reach any level, short perhaps of 1, without one having justification for believing P¹³.

¹² It is a consequence of this definition that E normically supports P iff E normically supports P more strongly than $\sim$ 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 $\sim$ P. More generally, given this definition of comparative normic support, it will turn out that, if there is a proposition Q such that E normically supports P more strongly than Q, then E must normically support P. Once again, the corresponding principle for justification seems questionable: If there is a proposition Q such that 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 automatically derivable from the former principles. I will have more to say about these issues in chapter 5.

¹³ 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. If the underlying probability space is infinite, then the fact that a proposition P has a probability of 1 given evidence E is consistent with there being possible worlds at which E holds and P does not. If these are amongst the most normal worlds at which E holds, E will not normically support P. Suppose my evidence is sufficient 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, much like this). For any integer n, this interval 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 sub-intervals as any other. The evidential probability that the speed of light is any particular value x within the interval must, then, be less than 1/n for any integer n – that is to say, 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 has an evidential probability of 1, this is perfectly consistent with there being possible worlds, consistent with my evidence, in which the speed of light is equal to x. If these worlds lie amongst the most normal

One consequence of choosing to model normic support in terms of possible worlds is that it will make the notion too coarse to discriminate between propositions that share the same possible worlds truth conditions. Propositions that are true at exactly the same worlds must, on the present model, have the same normic support profile – and, in particular, propositions that are true at all worlds will be normically supported by all bodies of evidence while propositions that are true at no worlds will be normically supported by none. One might regard this as an important objection to the present model – but to dismiss it on these grounds alone would, I think, be hasty.

One thing to note right away is that the notion of evidential probability, as standardly modelled, is *just as coarse*. Probability functions are standardly defined over sets of possible worlds – with the probability of a proposition taken to be determined by the probability of the set of worlds at which it's true (see, for instance, Williamson, 2000, section 10.4, Hájek, 2003, Douven and Williamson, 2006). As such, propositions that are true at exactly the same possible worlds will receive the same evidential probability – and, in particular, propositions that are true at all worlds will receive probability 1 relative to any possible body of evidence, while propositions that are true at no worlds will receive probability 0.

There are various ways of attempting to make evidential probabilities more discriminating – and these options are available in the case of normic support as well. I'll have more to say about this in section 6.2. In my view, though, issues of this general kind need to be understood as by-products of attempting to treat notions like normic support or evidential probability in a formal way. In building formal models of normic support or of evidential

worlds that are consistent with my evidence, then my evidence will not normically support the proposition that the speed of light is not equal to x .

probability, we strike a familiar bargain: We secure a kind of systematicity at the cost of a certain fineness of grain that may appear to be present in the informal notion.

There is, of course, no obligation to offer any formal model of normic support – we might work exclusively with the informal notion, defined in terms of explanation and the associated notion of normalcy. One of the benefits of a formal model, though, is that it allows us to definitively settle questions of logic. It will turn out, for instance, that normic support, as modelled here, is closed under multiple premise deductive consequence – if a body of evidence normically supports each of a set of propositions, then it will also normically support any proposition that they jointly entail. If evidence E normically supports a proposition P , then the most normal E -worlds are P -worlds. If evidence E normically supports a proposition 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 reasoning can be easily generalised for any number of premises.

It's worth noting that the informal characterisation of normic support already made it very plausible that it should possess this property. If the falsity of P would require special explanation given one's evidence and the falsity of Q would require special explanation given one's evidence, then presumably the falsity of $P \wedge Q$ would also require special explanation, given one's evidence. After all, the falsity of $P \wedge Q$ must involve either the falsity of P or the falsity of Q . The formal model, however, allows us to set this impression in stone.

Thus far, normic support has been proposed only as a necessary condition for justification. As such, the fact that normic support is closed under multiple premise deductive consequence won't guarantee that justification is similarly closed, though it will leave open the path to such a view. Whether we ultimately have this result will depend on what else, if anything, we take justification to require. I shall return to this issue in chapter 4.

The present model also predicts, relatedly, that no body of evidence could provide simultaneous normic support for each of an inconsistent set of propositions. If E normically supports P and normically supports Q, then P holds at the most normal worlds at which E holds and Q holds at the most normal worlds at which E holds. It follows that there are possible worlds at which P and Q both hold in which case P and Q are consistent. This reasoning can easily be generalised for any number of premises. Once again, this conclusion rings true even for the informal characterisation of normic support. If the falsity of P would require special explanation given one's evidence and the falsity of Q would require special explanation given one's evidence, then the falsity of $P \wedge Q$ would require special explanation given one's evidence. But the falsity of a contradiction never requires special explanation.

As discussed in chapter 1, a body of evidence *could* simultaneously make each of an inconsistent set of propositions arbitrarily likely – and it is for this reason that high evidential probability does not normatively coincide with knowledge. Perhaps one could never prove that normic support and knowledge normatively coincide without giving a substantial theory of knowledge – but there is no *formal* obstacle to their doing so. Unlike high evidential probability, normic support has the right logical features in order to normatively coincide with knowledge.

2.5. Objections and Replies

As I suggested above, my description of the laptop example and, in particular, of the evidence available to me, may be 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.

There is something dissatisfying about the way in which the knowledge account of evidence would have us treat this case – though it's difficult to put one's finger on exactly what it is. Part of what is puzzling about the case is that my belief in P_1 seems as though it could qualify as knowledge, while my belief in P_2 does not, even though my evidence for P_2 seems to be stronger than my evidence for P_1 . One way to solve this puzzle is to argue that my evidence for P_1 really is stronger than it appears. According to the knowledge account, P_1 is indeed more strongly supported by my evidence than P_2 – but this is only *because* my belief in P_1 is taken to qualify as knowledge. It's difficult to shake the impression that this is a kind of sleight of hand – and that *nothing* has been explained at all.

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 (unknown 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 not¹⁴.

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 at all inconsistent with this description. On the contrary, provided the details

¹⁴ According to epistemological disjunctivists, (McDowell, 1982, Pritchard, 2012) there is a fundamental epistemic asymmetry between veridical and non-veridical perceptions, such that the former put one in a kind of 'direct contact' with a fact in the world, while the latter do not. Disjunctivists may also take exception to my original description of the laptop example arguing that P_1 becomes a part of my evidence when I veridically perceive that Bruce's laptop is displaying a blue background. Once again, this objection can be answered by pointing out that it is in no way 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 is also important to point out that examples of this sort need not involve a contrast between statistical and perceptual evidence – as the Blue-Bus case illustrates, we can, for instance, build a corresponding example around a contrast between statistical and testimonial evidence.

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 the knowledge account of evidence might worry that I'm underestimating the evidence that's available to me, others may be concerned that I'm *overestimating* this evidence. Perhaps E_2 is not the kind of proposition that could literally be a part of my evidence. If my laptop is working and is properly implementing the algorithm, then it will be 99.9999% likely to be displaying a blue background – but I could perhaps have some doubts as to whether my laptop is working and is properly implementing the algorithm.

Let $\Pr$ be my evidential probability function at the point I set eyes on Bruce's laptop. Suppose that E_2 is not a part of my evidence but itself has some evidential probability that is close to but less than 1. This will make the evidential probability of P_2 more difficult to calculate, but it need not make any difference to the overall thrust of the example. If $0 < \Pr(E_2) < 1$ then, by the theorem of total probability, $\Pr(P_2) = \Pr(P_2 \mid E_2) \cdot \Pr(E_2) + \Pr(P_2 \mid \sim E_2) \cdot \Pr(\sim E_2)$. $\Pr(P_2 \mid E_2)$ is, of course, equal to 0.999999. $\Pr(P_2 \mid \sim E_2)$ will be more difficult to determine, but it will take some non-zero value. Let x be the value of $\Pr(P_1)$ – the probability, given my evidence, that Bruce's laptop is displaying a blue background. Provided, then, that $\Pr(E_2) > x/0.999999$, P_2 will turn out to be more likely, given my evidence, than P_1 – as the example requires. This result holds irrespective of the value of $\Pr(P_2 \mid \sim E_2)$.

More generally, if we let $\Pr(P_2 \mid \sim E_2) = y < 0.999999$, it is enough that $\Pr(E_2) > (x - y)/(0.999999 - y)$.

Proof

$$\Pr(E_2) > (x - y)/(0.999999 - y)$$

$$\Rightarrow \Pr(E_2) \cdot (0.999999 - y) > x - y$$

$$\Rightarrow \Pr(E_2).0.999999 - \Pr(E_2).y > x - y$$

$$\Rightarrow \Pr(E_2).0.999999 + y - \Pr(E_2).y > x$$

$$\Rightarrow \Pr(E_2).0.999999 + (1 - \Pr(E_2)).y > x$$

$$\Rightarrow \Pr(E_2).0.999999 + (\Pr(\sim E_2)).y > x$$

$$\Rightarrow \Pr(P_2 | E_2).\Pr(E_2) + \Pr(P_2 | \sim E_2).\Pr(\sim E_2) > \Pr(P_1)$$

$$\Rightarrow \Pr(P_2) > \Pr(P_1)$$

QED

If, for instance, $\Pr(P_1) = 0.95$ and $\Pr(P_2 | \sim E_2) = 0.4$ then, provided that $\Pr(E_2)$ is greater than 0.92, P_2 will be more likely, given my evidence, than P_1 . In order for this example to serve its purpose, then, it is not necessary that E_2 be a part of my evidence or even that it be certain given my evidence. Rather, it is enough that E_2 meet a given evidential probability threshold, which can be determined as a function of $\Pr(P_1)$ and $\Pr(P_2 | \sim E_2)$.

Before concluding, it's worth noting that there are further examples, of the same apparent structure as the laptop example, that don't seem to rest on any assumptions about a person's evidence. The following is adapted from an example given by Don Fallis (1997): Suppose I'm presented with two very large numbers n_1 and n_2 and asked to determine, within a certain time limit, whether or not they are prime. I know of one sure-fire method for detecting primeness – but the method is cumbersome and difficult. I apply this method to n_1 and, after carefully checking through all of my calculations, arrive at the verdict that it is indeed prime. I then notice that most of my time has elapsed so resort to a different method for n_2 . This method is much quicker and easier, but it is not sure-fire – it may yield the verdict that a number

is prime even if it is in fact composite, though it is very *unlikely* to do so. I apply the method to n_2 and arrive at the verdict that it is prime.

Let's fill in a few further details: It has been proven that, if a number n is composite, then the majority of the natural numbers between 1 and n will stand in a certain, easily detected, numerical relationship to n , whereas if n is prime, then none of the natural numbers between 1 and n will stand in this relation to n (Rabin, 1980). Suppose that, in the time I have left, I manage to test a sample of numbers less than n_2 and find that none have the relation in question to n_2 . If my sample is sufficiently large, then the chance of this result, given that n_2 really is composite, may be very low.

Under these circumstances, I would seem to have more justification for believing that n_1 is prime than I do for believing that n_2 is prime. Indeed, it seems I would not be justified in believing outright that n_2 is prime. A number of further observations accompany this: If I announce, once the time runs out, that both numbers are prime then, with respect to n_2 , what I seem to be doing is making an educated *guess* (which can of course be a perfectly reasonable thing to do in a test). When it comes to n_1 however it wouldn't be at all natural to describe my announcement as a 'guess', educated or otherwise. And, once I'm outside of the test context, it seems that I'm still perfectly entitled to announce that n_1 is prime, but I shouldn't simply announce the same thing about n_2 – the most that I should say, it seems, is that it is overwhelmingly likely to be prime.

As noted, though, the first method for detecting primeness is much more difficult to apply than the second and, in spite of my careful checking, the probability that I've made an error in applying the first method will be higher than the probability that I've made an error in applying the second method. Once these probabilities of error are factored in, the overall probability that the first number is prime could very well be lower than the overall probability

that the second number is prime. If pressed about my own fallibility I might even judge the second number to be more *likely* prime than the first, and I might be quite right to do so¹⁵. This judgment is not based however on any assumptions about what my *evidence* consists of – I’m simply estimating the error risks associated with these methods themselves along with the error risks associated with my attempting to implement these methods. It’s not at all clear what we should consider to be my relevant evidence in this case – arguably, the very notion of evidence seems of limited use in helping us to understand this sort of epistemic predicament.

In this chapter I’ve argued that the risk minimisation conception of justification is flawed and have begun to sketch an alternative picture. In section 2.1, 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 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 ways in which a proposition can be more than probable, though less than certain. I have outlined one such way here.

¹⁵ Fallis (1997) describes this kind of example in the course of arguing that there is no epistemic basis for the current mathematical practice of rejecting probabilistic methods (such as the method described here) as a means of establishing mathematical truths. Against the background of a risk minimisation picture, this conclusion may be correct – a probabilistic method with a low rate of implementation error can minimise risk more effectively than a deductive method with a high rate of implementation error. The notion of normic support could however offer one way of legitimating the current practice. I won’t pursue this further here. For discussion of Fallis’s position, and of this particular example, see Easwaran (2009).