

# What Else Justification Could Be<sup>1</sup>

Martin Smith

University of Glasgow

According to a captivating picture, epistemic justification is essentially a matter of epistemic or evidential *likelihood*. While certain problems for this view are well known, it is motivated by a very natural thought – if justification can fall short of epistemic certainty, then what else could it possibly *be*? In this paper I shall develop an alternative way of thinking about epistemic justification. On this conception, the difference between justification and likelihood turns out to be akin to the more widely recognised difference between *ceteris paribus* laws and brute statistical generalisations. I go on to discuss, in light of this suggestion, issues such as classical and lottery-driven scepticism as well as the lottery and preface paradoxes.

## I. RISK MINIMISATION

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

---

<sup>1</sup> I have benefited from discussions of this material with a number of people. Particular thanks to Philip Ebert, Daniele Sgaravatti, Stewart Shapiro, Ernest Sosa and Crispin Wright. Thanks also to an audience at the first Basic Knowledge Workshop at the University of St Andrews in November 2007.

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: 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. There are a number of apparent differences between these two cases. In particular, while my belief in this 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 standard 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 pervasive and quite natural 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 (at least at a first approximation) when the likelihood of P is sufficiently high and the risk of  $\sim$ P is correspondingly low. Call this the *risk minimisation* conception of justification.

This general picture seems to be taken for granted by a broad range of epistemologists – many of whom have little else in common (in addition to Hawthorne and Pritchard, see Russell, 1948, chap. VI, Chisholm, 1957, pp28, Goldman, 1986, section 5.5, Pryor, 2004, pp350-351, 2005). The acceptance of this picture tends to be more self-conscious in the so-called ‘formal epistemology’ tradition where it has been brought to the fore via the paradoxes of rational acceptability. A number of careful refinements of the basic risk minimisation picture have been proposed in the hope of circumventing the lottery and preface paradoxes (see for instance Lehrer, 1974, chap. 8, Pollock, 1990, pp80-81). I shall come to this in section III.

There is undoubtedly something very attractive about this way of thinking. If epistemic justification can fall short of epistemic certainty, then what else could it possibly *be* if not evidential probability or likelihood? If securing epistemic justification does not require the complete *elimination* of error-risk, then what else could it possibly require if not its minimisation? I hope to present a possible answer to these questions here – but first a little groundwork is required.

I assume here that evidence is propositional – that one’s body of evidence consists of a stock of propositions (Williamson, 2000, section 9.5). Though here is not the place to defend any particular account of when a proposition qualifies as part of one’s body of evidence, everything that I say can be reconciled both with Timothy Williamson’s ‘knowledge account’ of evidence (Williamson, 2000, chap. 9) as well as traditional empiricist accounts.

A proposition is likely, for one, to be true, just to the extent that it is made likely by one’s body of evidence. It is worth emphasising that the probability in question here is not objective frequency or propensity, though the evidence bearing upon a proposition can, naturally, concern objective frequencies or propensities. The probability in question is also quite independent of one’s actual estimation of how likely it is that a proposition is true. Evidential probability is sometimes analysed in terms of the probability estimate that an *ideally rational* being would make, given a body of evidence. This undoubtedly comes close to what is intended – but falls prey to the problems that standardly plague counterfactual analyses (Shope, 1978, Williamson, 2000, chap. 9). In any case, no precise analysis is necessary for present purposes. The notion of evidential probability is easily explained using examples (typically involving lotteries, dice or random processes of some kind).

The conception of justification that I will work with here is deliberately minimal. What I mean when I say that one has justification for believing P is simply that it would be epistemically or intellectually appropriate for one to believe or accept that P is true, given the evidence at one’s disposal. When I say that one has more justification for believing P than Q what I mean is that it would be more appropriate,

given one's evidence, for one to believe that P than it would be for one to believe that Q.

Justification, I take it, is what makes a belief a good candidate for knowledge. I'm happy to leave this suggestion somewhat imprecise for the time being, but the basic idea is this: My belief is justified just in case I have *done my epistemic bit* – the rest, as it were, is up to fate. My belief will qualify as knowledge provided that the world obliges or cooperates – but *I* am not required to do anything further<sup>2</sup>. The traditional analysis of knowledge can, I think, be seen as a somewhat clumsy way of trying to capture this basic idea. One lesson of the Gettier cases is that there are other ways, besides falsity, in which the world can fail to cooperate. While this connection between justification and knowledge will be helpful for what is to come, I won't rely upon it too heavily.

I am taking it for granted here that believing or accepting a proposition as true involves a commitment over and above merely regarding it as probable. In one sense this seems quite obvious – to accept that P is *probable* is precisely to adopt an attitude that does *not* risk error if P is false. To accept that P is probable is a way of 'hedging one's bets' – of avoiding an outright commitment to P. It seems as though I can regard it as *extremely* probable that, say, the number of stars in the universe is a composite number or that ticket #5472 will lose the lottery without actually *believing* or *accepting* either of these things. This view of acceptance is, in any case, a part of

---

<sup>2</sup> Here is one way of making this suggestion more precise, due to Bird (2007): If a subject S with mental states M forms a belief, then that belief is justified provided there is some possible world in which S, with the same mental states, forms a corresponding belief that amounts to knowledge. This, I think, comes close to being right. However, given S5, this proposal would appear to have the consequence that the justificatory status of a belief is *necessitated* by one's mental states and the way the belief is formed. There may be things to be said in favour of this contention – but it does seem to be an excessively strong rendering of the intuitive suggestion in the text.

the risk minimisation picture as I see it. 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).

Obviously, if the evidential probability of P given evidence E is very high, then any subject in possession of such evidence is epistemically entitled to offer a generous estimate of P's probability. What I hope to show, in the forthcoming sections, is that a body of evidence can, in some cases, be so constituted as to allow for a very generous estimate as to the probability of a proposition P, whilst mandating a *suspension of judgment* as to whether or not P is in fact true.

Before proceeding, it is worth mentioning that, provided evidential probabilities conform to the probability calculus, justification will not, on the risk minimisation picture, be closed under what is sometimes called multi-premise deductive consequence. That is to say, if one has justification for believing P and justification for believing Q, it will not follow automatically, on the risk minimisation conception, that one has justification for believing the deductive consequences of P and Q. Conjunction aggregates risk – even if the probabilities of P and of Q exceed the justification threshold, the probability of  $(P \wedge Q)$  need not. In terms of epistemic probability, the conclusion of a multi-premise deductive argument can be weaker than any of the premises.

This is undeniably a curious result. If multi-premise closure fails, then there must be possible situations in which it is intellectually appropriate to accept the premises of a deductive argument, accept its validity and yet remain agnostic about its conclusion. At this point I won't make too much of this, however. In fact, I shall remain officially neutral on the principle here, though I will have a little more to say about it in section III. The risk minimisation conception would appear to be consistent with the slightly weaker idea that justification is closed under *single* premise deductive consequence.

## II. NORMIC SUPPORT

Consider the following example, adapted from one given by Dana Nelkin (2000, pp388-389): Suppose that I have set up my computer 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 turn on my computer and then go into the next room to attend to something else.

In the meantime Bruce, who knows nothing about how my computer's background colour is determined, wanders into the computer room and *sees* that the computer is displaying a blue background. He comes to believe that it is. Let's suppose, for the time being, that my relevant evidence consists of the proposition that (E<sub>1</sub>) it is 99.9999% likely that the computer is displaying a blue background, while Bruce's relevant evidence consists of the proposition that (E<sub>2</sub>) the computer visually

appears to him to be displaying a blue background (I will consider other ways of describing the evidential situation in due course.)

Here are a few preliminary observations about this case: If I were to believe or accept that the computer is displaying a blue background before walking back into the computer room, 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 Bruce's belief this way. Second, Bruce's belief 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. My belief, on the other hand, would not constitute knowledge even if it happened to be true. If there were a power failure before I had a chance to look at the computer screen, I might well think to myself 'I guess I'll never *know* what colour the background really was'. But Bruce certainly wouldn't think this.

If someone were to ask Bruce 'what colour is the computer background?', he would be perfectly epistemically entitled to reply 'it's blue'. But if someone were to ask *me* the same question, 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. Bruce is not required to do any further investigation into the background colour – even though he easily *could* by, for instance, asking others to have a look. I, on the other hand, ought to do more investigation – by, for instance, going and having a look myself – before I rest on my laurels.

The implication of these considerations is clear enough: Bruce has *justification* for believing that the computer is displaying a blue background while I do not. In spite of this, the proposition that the computer is displaying a blue background is more likely for me to be true than it is for Bruce. While Bruce's evidence  $E_2$  does make it highly likely that the computer is displaying a blue background, it clearly does not *guarantee* that it is. After all, Bruce could be hallucinating, or he could have been struck by colour blindness, or there could be some coloured light shining on the screen etc. These are all rather unlikely – but presumably the likelihood, given Bruce's evidence  $E_2$ , that the computer is displaying a blue background would be nowhere near as high as 99.9999%. This, of course, is precisely how likely the proposition is, given my evidence  $E_1$ <sup>3</sup>. In believing that the computer is displaying a blue background, Bruce is running a *higher* risk of error than I am.

Here, then, are a cluster of intuitions that don't appear to fit in all that 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

---

<sup>3</sup> One may think of this as an instance of what of David Lewis (1980) called the 'principal principle'. According to the principal principle, in a nutshell, the epistemic probability of a proposition  $P$  conditional upon a body of evidence  $E$  containing the proposition that the objective probability of  $P$  at time  $t = x$ , will itself be  $x$  provided that  $E$  contains only admissible evidence – in particular, no relevant evidence pertaining to things that happened after  $t$ . Time  $t$ , in this case, could be thought of as the time at which I turn on my computer and leave the room.

According to one view, however, objective probability values must be extremal in a deterministic world (see Shaffer, 2007). If one took such a view, and the world were deterministic, then my evidence in the case described above could not concern objective probabilities – it would need to be reinterpreted as evidence about hypothetical long-run frequencies or some such.

prejudice. Perhaps this is the right thing to say – but I shall investigate a possible way of giving these intuitions more credit.

Let's reconsider the relationship between my evidence  $E_1$ , Bruce's evidence  $E_2$  and the proposition that (P) the computer is displaying a blue background. Clearly, neither  $E_1$  nor  $E_2$  *entails* that P is true. It would be perfectly possible for  $E_1$  and  $E_2$  to both be true while P is false. Notice, though, that if  $E_2$  is true and P is false, then this would appear to be a circumstance crying out for *explanation* of some kind. If Bruce visually perceives that the computer 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 Bruce is hallucinating, or has been struck by colour blindness, or there is funny lighting etc. It can't be that Bruce *just misperceives* – there has to be more to the story.

The circumstance in which  $E_2$  and P are both true, we might say, is *explanatorily privileged* over the circumstance in which  $E_2$  is true and P is false.  $E_1$  and P, however, do not appear to stand in this relationship. Although it would be very unlikely for  $E_1$  to be true and P false, this is not something that would require special explanation. All of the random numbers that might be generated by my computer 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 Bruce's belief that the computer 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 his error – either in terms of perceptual malfunction or disobliging features of the environment. If my belief that the computer is displaying a blue background turned out to be false, there need not be any available explanation for my error. The buck, as it were, may simply stop with me and the way that I chose to form my belief.

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 *contrastively* explain other arrival times in relation to this default (see Pietroski and Rey, 1995).

This may 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 notion of explanatory privilege, and indeed of an explanatory hierarchy, arises. The notion of explanatorily privileged conditions can be usefully compared to the familiar idea of an *idealised* or simplified model of a potentially complex actual phenomenon.

I will not attempt here to give a full, philosophically satisfactory account of explanation and idealised normalcy – in fact, everything that I say here will be compatible with a number of different ways of thinking about this. My aim is more modest – to present, in rough outline, an alternative picture or a *kind* of normalcy that stands apart from the statistical conception that can so easily entrance us.

Given my evidence  $E_1$ ,  $P$  would frequently be true. Given Bruce's evidence  $E_2$ ,  $P$  would *normally* be true. That is, given Bruce's evidence, if  $P$  turned out to be false, then this could be explained via some independent interfering factor. We might say, in this case, that  $E_2$  *normically supports*  $P$ . Given  $E_1$ , if  $P$  turned out to be false, then no such explanation need be available.  $E_1$  does not normically support  $P$ . The

distinction between the  $E_1$ -P relationship and the  $E_2$ -P relationship might fruitfully be compared to the distinction between statistical generalisations and normic or *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 usually B’, ‘As are typically B’ – and generics that are ‘unmarked’ – generics of the form ‘As are B’ (see, for instance, Leslie, 2008. I hope to pursue these comparisons elsewhere).

Let’s suppose that possible worlds can be ranked according to their comparative normalcy and say, slightly more formally, that a body of evidence E normically supports a proposition P just in case the most normal worlds in which E is true and P is false are less normal than the most normal worlds in which E is true and P is true. Further, a body of 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 P is false are less normal than the most normal worlds in which E is true and Q is false<sup>4</sup>.

---

<sup>4</sup> I am assuming, for simplicity, 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 each P-world, there is a more normal P-world. If we relax this assumption, we will 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. The characterisation of comparative normic support would need to be revised accordingly. I ignore this complication in the body text.

If we wish to think about normic support as a kind of intensional propositional operator, then we should, in addition, specify that it be vacuously satisfied in the case of an impossible antecedent – an antecedent that is true at no possible worlds. Amongst other things this will serve to ensure that normic support is weaker than classical entailment. Since we are only interested in certain instances of this relation, however, and this is not a feature that a possible corpus of evidence could have, this makes no difference for present purposes.

Here, then, is a support relation that evidence may bear to a proposition that does not simply reduce to probabilification. 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 attributable to *mitigating circumstances* – the error can be explained in terms of disobliging environmental conditions, or cognitive or perceptual malfunction or some such. Errors that do not fall into this category are, in effect, errors for which I am *intellectually responsible*. And if I leave myself open to such error, then I am intellectually negligent – even if my convictions turn out to be right.

What I propose is that, in order for one to have justification for believing 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. 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. The probability of P given E can reach any level (short, perhaps, of one) without this condition being met. Thus, the probability of P given E can reach any level (short, perhaps, of one) without one having justification for believing P<sup>5</sup>. To borrow a turn of phrase used by Pietroski and Rey (1995, pp84), the

---

<sup>5</sup> I am inclined to think that there will be possible cases in which one lacks justification even for propositions that have a probability of one, given one's evidence. As is often noted, zero probability does not imply impossibility or even falsity, at least where infinite probability spaces are concerned. Suppose X is a standard, continuous random variable, such that the probability of it assuming any one of the particular values over which it ranges is zero. Although the probability that  $X = x$  is zero and the probability that  $X \neq x$  is one, this is perfectly consistent with there being possible worlds, compatible with my evidence, in which X assumes the value x. In this case, my evidence need not entail the proposition that  $X \neq x$  and need not normically support it either.

notion of justification answers to the need to idealise in a complex world, not the need to describe a chancy one.

On my account, 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 evidence E, can exceed the probability of Q, given E, without one having more justification for believing P than Q<sup>6</sup>.

The property of being normically supported by a body of evidence would appear to be closed under multi-premise deductive consequence. That is, if E normically supports P and E normically supports Q then E normically supports  $P \wedge Q$ . This, at any rate, is a consequence of the way that normic support has been modelled here. Further, given the model that has been employed here, it can be shown that the *degree* of normic support that a conjunction enjoys, as measured by the abnormality of the most normal worlds in which one's evidence E is true and the conjunction false, will in fact be equal to the degree of normic support enjoyed by the least supported

---

<sup>6</sup> The following is suggested by this proposal: If one has more justification for believing P than Q, then one must have justification for believing P. If the most normal worlds in which one's evidence E is true and P is false are *less normal* than other worlds in which E is true, then the most normal worlds in which E is true must be worlds in which P is true (assuming, of course, that there are maximally normal worlds at which E is true). Some gradable adjectives have a corresponding feature – for instance if S is *wetter* than T, it would appear to follow that S must be *wet*. It certainly sounds very odd to say 'S is wetter than T, but S is not wet'. It does not sound at all odd, in contrast, to say something like 'S is *taller* than T, but S is not *tall*'. It does sound odd, to my ear at least, to say 'One has more justification for believing P than Q, but one does not have justification for believing P'.

conjunct. In terms of normic support, the conclusion of a multi-premise deductive inference will be as strong as the weakest premise<sup>7</sup>.

Provided we hold our standards of normalcy (our ranking of possible worlds) constant, the most normal worlds in which E is true and  $P \wedge Q$  is false will 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. Thus, the degree of normic support that E affords  $P \wedge Q$  will be equal to either the degree of normic support that E affords P or the degree of normic support that E affords Q – whichever is lower.

As I hinted above, my description of the computer example and, in particular, of the evidence available to Bruce and myself, is potentially contentious – in fact, it takes for granted a certain kind of traditional, empiricist conception of evidence. One who is impressed by the *knowledge* account of evidence might object to my description as follows: Once Bruce *sees* that the computer is displaying a blue background, he simply comes to *know* that it is. According to the knowledge account of evidence, this suffices for the proposition to qualify as part of Bruce's body of evidence. In this case, the probability, conditional upon Bruce's evidence, that the computer is displaying a blue background will in fact be one and, thus, *will* exceed the probability that the computer is displaying a blue background, conditional upon my evidence, contrary to what I have claimed.

---

<sup>7</sup> Given a certain assumption, possible worlds could be assigned *natural numerical* normalcy degrees – w is normal to degree one (maximally normal), x is normal to degree eleven etc. – and normic support strength could be gauged in an associated way. The required assumption is that we cannot have infinite ascending chains of increasingly normal worlds – that is, the normalcy ranks must be *well-ordered*. I remain neutral on this here. Even without this assumption we are left with comparative truths about world normalcy and, correspondingly, about normic support strength. I hope to discuss these issues elsewhere.

I think there is something rather unsatisfying about the description that the knowledge account offers in this case – though it is rather difficult to put one’s finger on precisely what it is. The question that the case naturally prompts is something like this: ‘How is it that Bruce’s belief qualifies as knowledge while mine never could?’ To say that the evidential probability of P is higher for Bruce than it is for me might constitute one kind of informative answer – but *not*, it seems, if it is motivated by the knowledge account of evidence. In this case, the only reason that the evidential probability of P is taken to be higher for Bruce than for me is *because* Bruce’s belief is taken to qualify as knowledge whereas mine is not – in which case it is clear that the original question has not been answered at all. It would appear, in fact, that the knowledge account will preclude any informative answer that involves a comparison between the evidence available to Bruce and the evidence available to me. Far from offering us any insight into the case, this would only seem to deepen our puzzlement.

We need not pursue this further, however. Fortunately, it is relatively easy to circumvent this kind of objection. We need only point out that the force of the computer example does not depend in any way upon the *actual background colour*. We can simply imagine that the background colour is red and Bruce *really is* hallucinating (unbeknownst to him of course). In this case, even the proponent of the knowledge account would, presumably, have to restrict Bruce’s relevant evidence to something along the lines of  $E_2$  (see Williamson, 2000, pp198). But this changes nothing. It would still be true that Bruce has justification for believing that the background is blue (while I do not).

The example that I have used in this section is undoubtedly rather contrived. While there is nothing particularly remarkable about Bruce's epistemic situation, the evidence that I possess – evidence, in effect, about the probabilistic features of a random process – is somewhat unusual. I think that evidence of this kind makes for the most compelling example of evidence that is *purely* probabilifying – that supports a proposition in a purely probabilistic fashion. While I do think that other, more commonplace, kinds of evidence share this feature, this is not something I will argue here. What is important for present purposes is that justification and evidential likelihood can come apart in principle, regardless of the extent to which they come apart in practice.

### III. SOME PUZZLES

Lottery cases famously raise a number of related epistemic puzzles. Perhaps the best known of these is the so-called 'lottery paradox' first described by Henry Kyburg (1961). The puzzle is essentially this: Consider a fair lottery with one thousand tickets and one guaranteed winner. It seems that one would be justified in accepting, for any given ticket, that it will lose. Given the set up, one also knows that some ticket will win. If justification is closed under multi-premise deductive consequence then one would be justified in accepting a contradiction – namely, that no ticket will win and that some ticket will win.

Kyburg's preferred and influential strategy for resolving the problem was to reject the principle of multi-premise closure. One general problem with this approach is that, even without multi-premise closure, one is forced to tolerate a conclusion that

is somewhat unpalatable – namely, that one can be justified in holding each of a corpus of beliefs known to be jointly inconsistent.

Many epistemologists, in any case, have attempted to hold on to multi-premise closure and to motivate instead the denial of the first premise – that one has justification for believing, of any given ticket, that it will lose. There has been a certain tradition of attempts to engineer this result, whilst leaving the broad risk minimisation conception of justification as intact as possible. I don't propose to review these attempts here. The following, which is the essence of the proposal made by Lehrer (1974, chap. 8) is representative: One has justification for accepting a proposition if it has a probability that exceeds the threshold and, in addition, exceeds the probability of any other proposition that is negatively relevant to it (where a proposition P is negatively relevant to Q just in case the evidential probability of Q is lower on the assumption that P than otherwise). There has also been a certain tradition of arguments exposing such attempted revisions as disastrous (Douven and Williamson, 2006 is perhaps the most recent and comprehensive example).

My inclination, of course, is also to reject the first premise – but not to try and reconcile its rejection within a risk minimisation framework. As Jonathan Vogel observes:

...although winning a lottery on a particular ticket is unlikely or improbable, it would not be *abnormal* in some intuitive sense, for the ticket one holds to turn out to be a winner.

(Vogel, 1990, pp16)

I have claimed that, in order for one to have justification for believing a proposition P, one's body of evidence must normically support P. One's only relevant evidence in the lottery case is the fact that the lottery is fair and that there are one thousand

tickets. This evidence does not normically support, for any particular ticket, the conclusion that it will lose – the one thousand possible outcomes will all be on an explanatory par. Amongst the most normal worlds in which the evidence holds true will be worlds in which each ticket wins. Thus, one lacks justification for believing, of any particular ticket, that it will lose.

One reason that people have traditionally been squeamish about denying that one can justifiably believe that a given ticket will lose a fair lottery is for fear of inviting radical scepticism. If we cannot be justified in believing, of a single ticket, that it will lose, then surely we cannot be justified in believing much at all. Richard Foley clearly expresses this worry:

...most of us, on reflection, would be likely to regard the evidence we have for the claim that ticket one will not win to be as strong as the evidence we have for most of the claims that we believe. For instance, we would be likely to regard the evidence for it as being at least as strong as the evidence we have for the claim that the room we have just left still has furniture in it...

(Foley, 1987, pp245)

My response to Foley's concern is, perhaps, easy to anticipate: 'Strong', in this context, is effectively ambiguous. On one reading what Foley writes is true, on another it is false.

Suppose that one's relevant body of evidence E consists of the proposition that I saw furniture in the room 30 seconds ago and the proposition that I hold a single ticket in a fair lottery of one million tickets. Let P be the proposition that the room has furniture in it now and Q be the proposition that I will lose the lottery. What seems clear is that P is no more probable, given E, than Q is – indeed Q may be more probable. Foley supposes, however, that this is all that there is to say about the

bearing of the evidence – and this is a mistake. E normically supports P but does not normically support Q. If E were true and P false, there would need to be some explanation for this – such as very fast-working burglars or the interventions of a mischievous demon. While E *probabilistically supports* Q more strongly than P, it normically supports P more strongly than Q.

Even if this immediate threat is defused, though, a further, and somewhat subtler, sceptical menace still looms. As John Hawthorne (2003) writes:

Just as I have excellent statistical grounds for supposing that any given lottery ticket will lose, I have excellent statistical grounds for supposing that a given apparently healthy person will not have a fatal heart attack very soon. Just as there was no special reason in advance for supposing that the winning ticket was going to win, there was no reason in advance for expecting the worst for some heart attack victim who was apparently healthy. And just as many of our ordinary commitments entail that this or that person will lose a lottery, many of our ordinary commitments entail that this or that person will not soon suffer a fatal heart attack.

(Hawthorne, 2003, pp3)

If one cannot have justification for believing, in the absence of special, non-statistical reasons, that one will lose a fair lottery, then it seems as though one cannot have justification for believing, in the absence of special, non-statistical reasons, a range of relevantly similar things – for instance, that one will not soon have a fatal heart attack. Given that we often believe things that entail such propositions then, provided justification is closed under single premise deductive consequence, this would appear to usher in a widespread scepticism.

Although we have excellent statistical grounds for thinking that an apparently healthy person will not soon suffer a fatal heart attack, this is not to say however that we conceive of this possibility as the outcome of a completely *random* process, like a lottery. While there may be a kind of pessimistic mind-set in which we view future health as tethered to pure chance, generally speaking we regard the continued health of an apparently healthy individual as being more *normal*, in the sense that his suddenly suffering a fatal heart attack would be a circumstance requiring special explanation. That is, if an apparently healthy person were to suddenly drop dead of a heart attack, this would require special explanation in a way that his continued health would not – so, at any rate, we are inclined to suppose.

This is a significant disanalogy with propositions about lottery outcomes. Recall Vogel's claim above: It would not be abnormal, in some intuitive sense, for my ticket to win a fair lottery, no matter how unlikely that outcome may be. And yet we *would* regard it as abnormal, in just this sense, for an otherwise healthy individual to drop dead of a heart attack. On the view I have developed, one can be epistemically entitled to disregard possibilities that are genuinely abnormal – but not possibilities that are merely *unlikely*<sup>8</sup>.

Similar remarks, I think, can apply to Hawthorne's other examples of so-called 'lottery propositions' such as *the president has not died within the last few minutes*, *my refrigerator is still running* and *my plane will not crash*. The analogy between

---

<sup>8</sup> If I have justification for believing, of an apparently healthy individual, that he won't soon suffer a fatal heart attack then, given the multi-premise closure principle, I can presumably have justification for believing, for arbitrarily large collections of apparently healthy people, that none will soon suffer fatal heart attacks. This result does not strike me as clearly disastrous – but it is for reasons of this kind that I remain neutral on multi-premise closure here. I will have a more to say about this in due course.

these propositions and propositions about lottery outcomes is not robust. Scepticism about the latter need not overflow into scepticism about the former.

While the class of ‘lottery propositions’ may not be as broad as Hawthorne suggests, it may of course turn out that some of our ordinary commitments do entail the rejection of possibilities that are merely unlikely, as opposed to genuinely abnormal – certainly there can be no general a priori guarantee to the contrary. In my view, we are not justified in undertaking any such commitments – at least not without some probabilistic qualification. The question as to which commitments may fall into this class is a somewhat delicate one, and one I don’t propose to investigate further here.

Another puzzle associated with lottery intuitions is that emphasised by Harman (1968), Cohen (1988, section IV) and DeRose (1996) (see also Sorensen, 2006): Suppose that Bob has bought a single ticket in two lotteries, both of which are guaranteed to have a single winner. The first of these is a relatively small community lottery in which just one hundred tickets were sold and the other is a nationwide lottery with a total of one million tickets. Suppose that both of these lotteries have now been drawn and Bob has not won either. Bob watches the drawing of the community lottery on TV and hears the announcer read out the number of the winning ticket which, naturally, is not his. The broadcast of the nationwide lottery draw, however, has been delayed and Bob has not yet seen it.

It seems clear, in this situation, that Bob *knows* that he has lost the community lottery, but does not yet know that he has lost the nationwide lottery. A number of

further intuitions seem to accompany this: It seems reasonable for Bob to assert that he has lost the community lottery, but premature for him to assert that he has lost the nationwide lottery. If asked about the nationwide lottery, Bob ought to say something like 'It is *very likely* that I've lost' or 'I've *almost certainly* lost'. This sort of qualification is not required when it comes to the community lottery. It also seems reasonable for Bob to throw away his community lottery ticket stub or to sell it for a cent (perhaps to someone who has not yet seen the draw). But it would certainly not be reasonable for Bob to throw out his *nationwide* lottery ticket stub or to sell that for a cent.

In spite of all of this, it is very plausible to think that the probability, given Bob's evidence, that he has lost the community lottery is actually *lower* than the probability, given that same evidence, that he has lost the nationwide lottery. After all, what Bob sees on TV does not guarantee that he has lost the community lottery. The person drawing the lottery could have misread the ticket, or hallucinated a different number etc. Such things may not happen often but presumably the likelihood of such an occurrence would be somewhat higher than one in ten thousand. How is it that Bob can know that he has lost the community lottery, reasonably assert that he has lost the community lottery and even *reasonably tear up his ticket stub* when the likelihood, for him, that he has lost the community lottery is actually lower than the likelihood that he has lost the nationwide lottery?

Perhaps the most common putative solution to this problem exploits the notion of *sensitivity* (Dretske, 1971, DeRose, 1996). If Bob had won the community lottery then he *would* have heard his ticket number read out on TV. That is, if Bob had won

the community lottery, his evidence would have been different. We might say, in this case, that Bob's evidence is *sensitive* to his losing the community lottery. In contrast, if Bob had won the nationwide lottery, he wouldn't have seen or heard anything different – his evidence would have been exactly the same.

This is a striking observation – but it cannot be the full story as to what is driving our intuitions. Imagine now a slightly different case in which Bob *has* won the community lottery, but the person drawing the lottery actually *does* hallucinate a completely different number on the winning ticket. As a result, he reads out the wrong number at the televised draw that Bob sees. It is *just* as reasonable, in this case, for Bob to assert that he has lost the community lottery. It is just as reasonable for Bob to throw away his community lottery ticket stub or to sell it for a cent. As tragic as such actions would be, we could I think *forgive* Bob for taking them. And if he did act in this way, he may well have a legitimate claim to compensation of some kind. Bob, of course, does not *know* in this case that he has lost the community lottery, but his belief would appear to be a very good candidate for knowledge.

The notion of sensitivity, needless to say, is no help here. Bob's evidence, in this case, is clearly *not* sensitive to his having lost the community lottery – after all he has actually *won* the community lottery. Of course, if we are able to resolve the puzzle in this more difficult case, then presumably we won't need to invoke sensitivity considerations in order to resolve it in the former case.

This new case also shows us that the knowledge account of evidence cannot offer a satisfactory resolution of this puzzle. With respect to the original case, the

proponent of the knowledge account could protest that the proposition that Bob has lost the community lottery is part of Bob's evidence and thus more likely, given that evidence, than the proposition that he has lost the nationwide lottery. But this cannot be maintained with respect to the new case. In the new case, the proponent of the knowledge account will have to admit that Bob's having lost the nationwide lottery is more likely, given his evidence, than his having lost the community lottery.

My account of justification offers a straightforward resolution that will apply in both cases: If Bob had won the community lottery in spite of what he saw on TV, then this would require some explanation – the announcer hallucinated or misread the ticket number or some such. In contrast, if Bob had won the nationwide lottery, in spite of the odds against it, then no special explanation would be needed. Someone's ticket had to win – it might just as well have been Bob's as any other. The one million possible outcomes are on an explanatory par. Bob's evidence (in both the easy and difficult case) normically supports the conclusion that he has lost the community lottery, but does not normically support the conclusion that he has lost the nationwide lottery. As a result, Bob has justification for believing that he has lost the community lottery, but not for believing that he has lost the nationwide lottery. This is why it is reasonable for Bob to assert that he has lost the community lottery, but not that he has lost the nationwide lottery. This is why it is reasonable for Bob to act on the proposition that he has lost the community lottery, but not on the proposition that he has lost the nationwide lottery.

The fact that one should not in general assert that one has lost a fair lottery, purely on the basis of the odds involved, has been observed by others – notably

Williamson (2000, section 11.2) and Hawthorne (2003, section 1.3). Both Williamson and Hawthorne take this observation to support the *knowledge* account of assertion, according to which one should only assert what one *knows* to be true. Put simply, their reasoning is something like this: If one should not assert that one has lost a fair lottery, then no epistemic probability threshold shy of one can suffice for warranted assertion. Since only one's knowledge has an epistemic probability of one, one should only assert what one knows to be true.

The assumption here is clear enough: Any appropriate epistemic standard less demanding than knowledge will have to be understood in terms of an epistemic probability threshold, shy of one. Williamson does consider the possibility that evidence may, in some cases, support propositions in a way that does not reduce to probabilification, but then falls back on something like the risk minimisation conception, claiming that, if one has warrant to assert a proposition, then surely one has warrant to assert any proposition that is *more likely* for one to be true (Williamson, 2000, pp251). Once we give up this principle, we can see that the knowledge account of assertion draws little support from consideration of lottery cases (see Nelkin, 2000, Douven, 2006, pp464, Kvanvig, 2007).

A third epistemic puzzle associated with lotteries is that identified by Jonathan Vogel (1999) and discussed by Hawthorne (2003, section 1.2): Suppose that every year Sue's community hosts a hundred ticket lottery and Sue is in the habit of buying a ticket. Clearly, I cannot know, purely on statistical grounds, that Sue will not win this year's community lottery. However, it seems that I *can* know that Sue is not going to win the community lottery *every year for the next fifty* – at least this strikes

many as intuitive. But my only evidence for believing this, presumably, is statistical. Furthermore, if Sue were to buy a single ticket in a mammoth lottery with  $10^{100}$  tickets, we would once again have the intuition that I cannot know that she'll lose – even though, given a few idealising assumptions, the odds of this happening are exactly the same.

Notice, though, that Sue's winning the community lottery fifty years running would be a circumstance that would undoubtedly rouse our *suspicious*. More precisely, it would be a circumstance that would motivate us to seek out some *special explanation* – the lottery is rigged or biased or a guardian angel is smiling on Sue or some such. We can easily imagine ourselves exclaiming 'Sue has won the lottery every year for the past fifty?! There has to be some explanation!' What this shows is that we are not inclined to treat the  $10^{100}$  possible outcomes of fifty community lotteries as being on an explanatory par – although we find it natural to treat the  $10^{100}$  mammoth lottery outcomes this way. In so far as one thinks that Sue's winning the community lottery fifty years running would be a circumstance demanding special explanation, one will regard oneself as having tangible reason to prefer the other possible outcomes.

My only evidence for thinking that Sue will not win fifty years running is probabilistic – and yet it is extremely tempting for me to think about the situation in *non-probabilistic* terms. There is considerable evidence to suggest that people are strongly disinclined to accept that event sequences exhibiting a recognisable *pattern* could arise through a purely random process. People tend to perceive such sequences as nonchance and as demanding special explanation (see, for instance, Falk and

Konold, 1997). It is often noted anecdotally that humans have a general tendency to disregard the possibility of ‘coincidental’ events occurring purely by chance – preferring, in many cases, even conspiratorial or supernatural explanations over such a concession. Some research tends to corroborate this impression<sup>9</sup>. The urge to say that I can know in advance that Sue will not win the community lottery fifty years running is, I suggest, an urge that springs from essentially the same tendency.

It is our bias against patterned outcomes that discourages us, in the iterated lottery situation, from seeing Sue’s winning fifty years running as being on a par with the other possible outcomes. No such bias will influence our reasoning about the mammoth lottery. Further, it seems possible to divest oneself of the troublesome intuitions – the comparison between the iterated lottery and the mammoth lottery and the associated parity of reasoning considerations can, I think, be helpful here. Once I do genuinely conceptualise Sue’s winning fifty years running as just one amongst  $10^{100}$  possible winning ticket combinations, each equally probable and equally plausible, then I will lose any tendency to think that I can simply rule it out.

---

<sup>9</sup> There is some evidence indicating that people tend to perform badly when asked to generate random sequences or to predict the outcomes of random processes – in particular, people appear to exhibit a bias against sequences with a large amount of repetition and tend to overestimate the frequency of alternation (see, for instance, Ross and Levy, 1958). People’s *judgments* regarding the randomness of presented sequences have also been the focus of research (Falk, 1981, Lopes and Odin, 1987, Falk and Konold, 1997). Findings tend to corroborate the existence of a general bias to associate randomness with excessive alternation. Naturally, the interpretation of the experimental findings is delicate and has been subject to dispute – for an overview see Bar-Hillel and Wagenaar (1991), Nickerson (2002).

One study, performed by Falk and Konold (1997), demonstrated that the average confidence with which subjects judged sequences of binary values to be randomly generated was strongly correlated with the average *time* it took for subjects to successfully memorise such sequences and that both judgments of randomness and memorisation performance were well predicted by a measure of encoding difficulty. Any general hypothesis to the effect that the ease of digesting or encoding a sequence of events guides our judgments of randomness would go some way toward explaining a strong verdict of nonrandomness in the case of a repetitive sequence such as Sue’s winning the community lottery fifty years running.

I have argued that one can lack justification for believing propositions that are nevertheless highly likely for one to be true. In order for one to have justification for believing that proposition P is true, P must be normically supported by one's body of evidence, and evidence can fail to normically support propositions that are nevertheless highly likely conditional upon that same evidence. This naturally invites the question of whether we ought to admit cases of the opposite kind – that is, cases in which one *has* justification for believing propositions that are nevertheless highly *unlikely* for one to be true. I shall conclude by considering, in a rather more speculative spirit, whether we should.

It is clear that a body of evidence can normically support propositions that are nevertheless highly unlikely conditional upon that same evidence. This will follow straightforwardly, given that normic support is closed under potentially risk-aggregating multi-premise deductive consequence relations. Familiar preface paradox type cases provide clear examples of this phenomenon: Suppose I am organising a dinner party to which I have invited fifty guests. All have responded saying that they will attend. Plausibly, I have evidence normically supporting the propositions that Chuck will attend, that Barbara will attend etc. Presumably, these propositions are individually quite likely for me to be true – let's say for the sake of argument that each is 90% likely.

Given that normic support is closed under multi-premise deductive consequence, I have evidence normically supporting the proposition that *every* guest will attend. If this proposition turns out to be false, then there must be some guest who fails to attend and, given his positive reply, some explanation as to why this is –

he lied or caught a cold or got into a car accident or some such. But, in spite of this, the proposition that every guest will attend is very unlikely, given my evidence, to be true – indeed, it is far more likely to be false. Assuming that the probabilities of each guest attending are independent, the probability of *every* guest attending will be in the region of one in two hundred.

Do I, then, have justification for believing that every guest will attend? Nothing I have said so far commits me to a positive answer – normic support has been proposed as a *necessary* but not a sufficient condition for justification. Indeed, the view I have developed is quite compatible with the claim that there is some evidential probability requirement upon justification. According to one sort of view, although justification will not exclusively be a matter of risk minimisation, it will have a risk minimisation *component*. Multiple premise closure for justification will, on this view, have to be sacrificed – and the risk-aggregating ‘preface paradox’ cases will be the very cases in which the principle fails. Though it lacks a certain overall elegance, this ‘hybrid’ view is in many ways quite attractive. I do think however that it is worth exploring a more radical alternative – a view on which multiple premise closure is maintained and justification is deemed compatible, in principle, with evidential probabilities that are arbitrarily low.

On this view, justification will have no essential connection to the minimisation of error risk. This is not to say, of course, that securing justification will not require the successful *management* of error risk – it is just that managing error risk and minimising error risk need not be the same thing. There is a sense in which both probabilistic support and normic support should be understood as providing a *kind* of

protection against error. Probabilistic support makes for what we might call a ‘low-risk’ guarantee – it serves to minimise the *overall* error risk. This is why the guarantee can evaporate as propositions are conjoined and risk aggregates. Normic support is not like this. It makes for a ‘no risk’ guarantee – that is, it serves to *eliminate* error risk of a *certain kind* (roughly speaking, the risk of error that is not attributable to, or explicable in terms of, extenuating circumstances). If each of a set of propositions enjoys normic support then, in conjoining them, there is *no risk* of the relevant kind to aggregate and, thus, the sort of guarantee offered cannot weaken.

It is often observed just how odd it would seem to accept a series of propositions, accept that further propositions follow deductively and yet remain agnostic about those further propositions. It would be very strange for me to accept, in a single conversational setting, that Chuck will attend, Barbara will attend, Pam will attend, running right through my invitation list, and then balk at the conjunctive claim that everyone will attend – without withdrawing or weakening any of my previous commitments. Further, in those circumstances in which I would be most obviously anxious about accepting the conjunction – when the possibility of car accidents or lies or the poor statistical track record of such arguments is salient to me – I would also be somewhat anxious about accepting flat-out the individual conjuncts, preferring something probabilistically qualified. With such concerns at the forefront of my mind, I’d find myself inclined to say ‘Chuck will most probably attend’, ‘Barbara will very likely attend’ etc.

The claim that it can be intellectually appropriate for me to believe that every guest will attend does undeniably make us uneasy. But the alternative is not clearly

preferable. For the alternative, after all, is that I am intellectually *mandated* to deny or suspend judgment on this proposition even while it is perfectly appropriate for me to accept, for every individual guest, that he or she will attend. It seems that no one would self consciously hold to such a combination of attitudes – and we should be somewhat suspicious of any theory that would make it intellectually appropriate to do so.

In this paper I have attempted to disarm the natural thought that lies behind the risk minimisation conception of justification. I have attempted to outline a kind of positive epistemic status that is defeasible, falls shy of knowledge and of epistemic certainty, but is nevertheless quite independent of epistemic probability or likelihood. There are different ways of responding to human fallibility – for managing the ever-present possibility of error. In the end, perhaps the only real failing of the risk minimisation picture is a certain myopia in this respect.

### References

- Adler, J. (2002) *Belief's Own Ethics* (Cambridge, MA: MIT Press)
- Bar-Hillel, M. and Wagenaar, W. (1991) 'The perception of randomness' in *Advances in Applied Mathematics* v12, pp428-454
- Butler, J. [1736] *The Analogy of Religion* (London: J.M.Dent and Co., 1906)
- Carnap, R. (1950) *The Logical Foundations of Probability* (Chicago: University of Chicago Press)
- Cavell, S. (1979) *The Claim of Reason* (Oxford: Oxford University Press)
- Chisholm, R. (1957) *Perceiving* (Ithaca, NY: Cornell University Press)

- Christensen, D. (1999) 'Measuring confirmation' *Journal of Philosophy* v96, pp437-461
- Cohen, S. (1988) 'How to be a fallibilist' *Philosophical Perspectives* v2, pp91-123
- 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)
- Douven, I. (2006) 'Assertion, knowledge and rational credibility' *Philosophical Review* v115(4)
- Douven, I. and Williamson, T. (2006) 'Generalising the lottery paradox' *British Journal for the Philosophy of Science* v57(4), pp755-779
- Diener, D. and Thompson, W. (1985) 'Recognizing randomness' *American Journal of Psychology* v98 pp433-447
- Dudman, V. (1992) 'Probability and assertion' *Analysis* v52, pp204-211
- Eells, E. and Fitelson, B. (2000) 'Measuring confirmation and evidence' *Journal of Philosophy* v97, pp663-672
- Falk, R. (1981) 'The perception of randomness' *Proceedings of the Fifth International Conference for the Psychology of Mathematics Education* (Grenoble: Laboratoire IMAG)
- Falk, R. and Konold, C. (1997) 'Making sense of randomness: Implicit encoding as a basis for judgment' *Psychological Review* v104, pp301-318
- Field, H. (2003) 'No fact of the matter' *Australasian Journal of Philosophy*, v81, pp457-480
- 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)
- 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)
- Goldman, A. (1986) *Epistemology and Cognition* (Cambridge, MA: Harvard University Press)
- Harman, G. (1986) *Change in View* (Cambridge, MA: MIT Press)
- Hawthorne, James (1996) 'On the logic of nonmonotonic conditionals and conditional probabilities' *Journal of Philosophical Logic* v25, pp185-218
- Hawthorne, James and Bovens, L. (1999) 'The preface, the lottery and the logic of belief' *Mind* v108, pp241-264
- Hawthorne, John. (2003) *Knowledge and Lotteries* (Oxford: Oxford University Press)
- Howson, C. and Urbach, P. (1993) *Scientific Reasoning: The Bayesian Approach* (Chicago: Open Court)
- Hume, D. [1739] *A Treatise of Human Nature* (London: Dent, 1911)
- 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)
- Kaplan, M. (1981) 'A Bayesian theory of rational acceptance' *Journal of Philosophy* v78, pp305-330

- Kaplan, M. (1996) *Decision Theory as Philosophy* (Cambridge: Cambridge University Press)
- 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)
- Lehrer, K. (1974) *Knowledge* (Oxford: Clarendon Press)
- Leslie, S. (2008) 'Generics: Cognition and acquisition' *Philosophical Review* v117(1), pp1-47
- Lewis, D. (1980) 'A subjectivist's guide to objective chance' in *Studies in Inductive Logic and Probability* vol. II (University of California Press)
- Locke, J. [1689] *An Essay Concerning Human Understanding* (Oxford: Clarendon Press, 1975)
- Lopes, L. and Oden, G. (1987) 'Distinguishing between random and non-random events' *Journal of Experimental Psychology: Learning, Memory and Cognition* v8, pp626-636
- Maher, P. (1993) *Betting on Theories* (Cambridge: Cambridge University Press)
- Makinson, D. (1965) 'The paradox of the preface', *Analysis* v25, pp205-7
- 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
- Nickerson, R. (2002) 'The production and perception of randomness' *Psychological Review* v109 pp330-357

- Okasha, S. (2004) 'Wright on the transmission of support: A Bayesian analysis' *Analysis* v64, pp139-146
- 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)
- 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)
- Ross, B. and Levy, N. (1958) 'Patterned prediction of chance events by children and adults' *Psychological Reports* v4, pp87-124
- 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)
- Schaffer, J. (2007) 'Deterministic chance?' *British Journal for the Philosophy of Science* vv58, pp113-140
- Schiffer, S. (2003) *The Things We Mean* (Oxford: Oxford University Press)
- Shope, R. (1978) 'The conditional fallacy in modern philosophy' *Journal of Philosophy* v75, pp397-413

- Smith, M. (2007) 'Ceteris paribus conditionals and comparative normalcy', *Journal of Philosophical Logic*, v36(1)
- Smith, M. (2009) 'Transmission failure explained' *Philosophy and Phenomenological Research* v79(1)
- Sorensen, Roy, (2006) 'Epistemic paradoxes', *The Stanford Encyclopedia of Philosophy (Fall 2006 Edition)*, Zalta, E. ed.  
 URL = <<http://plato.stanford.edu/archives/fall2006/entries/epistemic-paradoxes/>>
- Stanley, J. (2005) *Knowledge and Practical Interests* (Oxford: Clarendon Press)
- Sturgeon, S. (2008) 'Reason and the grain of belief' *Nous* v42(1)
- 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
- Williamson, T. (2000) *Knowledge and its Limits* (Oxford: Oxford University Press)