

Epistemic Entitlement, Leaching and Epistemic Risk

Luca Moretti and Crispin Wright

ABSTRACT

According to Crispin Wright, we have evidential justification for, or knowledge of, various propositions that we quotidianly accept only if we have antecedent justification for accepting general hinge propositions—called ‘cornerstones’—which cannot be evidentially supported. Wright contends that this doesn’t elicit scepticism, for we are non-evidentially entitled to accept cornerstones. This paper focuses on the Leaching Worry—the concern that since the epistemic risk of accepting a cornerstone *C* without evidence for it is significantly high, the epistemic risk of accepting a proposition *P* for which *C* is a cornerstone is also significantly high, to the effect that one cannot have evidential justification for, or knowledge that, *P*. We suggest that Wright’s original response to the Leaching Worry retains its strength if risk is construed in accordance with two non-orthodox (non-probabilistic) notions of risk recently introduced by Duncan Pritchard (modal notion) and Philip Ebert, Martin Smith and Ian Durbach (normic notion). We concede, however, that Luca Moretti’s recent probabilistic variant of the Leaching Worry isn’t undercut by Wright’s original response. We put forward two novel responses: revising the notion of significant epistemic risk originally adopted by Wright, or broadening the range of attitudes that we are non-evidentially entitled to.

1. Introduction

One type of argument to sceptical paradox proceeds by making a case that a certain kind of metaphysically “heavyweight” (Dretske, 2005, Hawthorne 2005) or “cornerstone” (Wright 2004) proposition is beyond all possible evidence and hence may not be known or justifiably believed. Examples of such propositions include that there is an external material world at all, that there are any other minds besides one’s own, that the world is replete with general lawlike patterns that are reliably reflected in local samples, and that the world did not come into being twenty minutes ago brimming with apparent traces (including apparent memories) of a much more substantial history. If such propositions are false, then sense experience is no guide to material reality, behaviour and physical condition is no guide to others’ mental states, generalisations inductively suggested by local samples are unreliable and most apparent evidence for the past has no probative force

The details of such sceptical arguments show some variety, but they do not matter for present purposes.¹ We are concerned here rather with a general strategy of response to them which argues that, even were the point admitted that the heavyweights and cornerstones are indeed beyond all evidence and knowledge, it would not follow that we thereby forfeit all claim to know, or at least justifiably to believe, the vast legion of specific propositions about the material world, others’ minds, scientific law and the past, that we are accustomed to think

¹ For one potentially important distinction, between what are there dubbed respectively “Humean” and “Cartesian” forms of sceptical paradox, see Wright (2004).

that we can and do know, or at least justifiably believe. This strategy, defended by Wright,² pivots on a notion of *epistemic entitlement*, in accordance with which one may be rationally fully warranted in trusting in the truth of certain propositions even though they lie beyond the range of possible knowledge or evidence. The basic idea is then that, with rational trust in such entitlements in place, a subject can justifiably proceed to attach the evidential significance to the deliverances of their senses, other's behaviour and physical condition, local patterns in samples, and their apparent memories and other presumed evidential traces of the past, that they customarily do. So if the sceptical challenge is to demonstrate how it can be rational to take ourselves to know or anyway to have ample evidential justification for the ordinary quotidian empirical beliefs that we routinely incessantly form, the answer is that, once it is acknowledged that it is fully rational trustingly to accept the cornerstones for a domain of enquiry, there is no straightforward passage from the thesis that we have no evidence for them to the conclusion that we thereby forfeit any positive epistemic standing for our beliefs throughout the domain in question. Rather, if we are rationally entitled to trust that there is an external material world, e.g., then there is no obstacle to the thought that, at least as far as that particular question is concerned, our ordinary perceptual faculties are reliably delivering veridical information about it.

Again, our present concern is not with the case for thinking that we do indeed possess such epistemic entitlements, but with a train of thought that argues that they cannot actually do the counter-sceptical epistemic work just adumbrated. That is the gist of the Leaching Worry.

2. The Leaching Worry

Wright (2004) characterised the worry thus:

The general picture is that the cornerstones which sceptical doubt assails are to be held in place as things one may warrantably trust without evidence. Thus at the foundation of all our cognitive procedures lie things we merely implicitly trust and take for granted, even though their being entitlements ensures that it is not irrational to do so. But in that case, what prevents this 'merely taken for granted' character from leaching upwards from the foundations, as it were like rising damp, to contaminate the products of genuine cognitive investigation? If a cognitively earned warrant—say my visual warrant for thinking that there is a human hand in front of my face right now—is achieved subject to a mere entitled acceptance that there is a material world at all, then why am I not likewise merely entitled to accept that there is a hand in front of my face, rather than knowing or fully justifiably believing that there is? ³

One natural way to elaborate this doubt is in terms of a notion of *significant epistemic risk*, whereby one is at significant epistemic risk in accepting *P* just if one lacks both knowledge of *P* and a balance of evidence in favour of believing *P*. Conversely if one knows *P* or one's evidence on balance favours believing *P*, then doing so involves, in this specific sense, no significant epistemic risk. Let *C* be the cornerstone proposition that there is an external material world and *P* some specific material world proposition which one would normally regard as known or at least as comfortably justified on the basis of one's current perceptual experience. *P* will entail *C* and, although rationally entitled to do so, we have, according to the sceptical paradox, absolutely no evidence for *C* and are therefore at significant epistemic risk, as defined above, in trusting in its truth. But were *C* not true, neither would *P* be. So how can we *not* be at significant epistemic risk in accepting *P*? And how can that admission be consistent with our having knowledge of, or a balance of evidence in favour of *P*?

² In addition to his (2004) see also Wright (2012) and (2014).

³ Wright (2004), p. 207.

As Sebastiano Moruzzi has observed,⁴ we can develop this Leaching worry in terms of the putative inconsistency of the three propositions, (1), (2) and (4), below:

(1) *If we run a significant risk in accepting C, then we run a significant risk in accepting P*

(1) may seem merely to articulate an immediate implication of C's being unevidenced and a logical consequence of P.

(2) *We do run a significant risk in accepting C*

This is entailed by the definition of “significant epistemic risk”. After all, entitled to accept C though we may be, the dialectical context is one where it has been conceded to the sceptic that we have no evidence for C.

Hence

(3) *We do run a significant epistemic risk in accepting P*

Yet if the idea that we are rationally entitled to trust in C is to do any interesting counter-sceptical work, it must somehow put us in position to claim that

(4) *We know P.*

And surely the propositions (4), that P is known, and (3), that we run a significant epistemic risk in accepting it, should be regarded as non-cotenable.

Wright (2004) concurred and responded that, since (2) is acknowledged on all hands, — and if deployment of the notion of entitlement is indeed to be of any counter-sceptical use, — we need to reject (1). But that, it was there contended, is exactly what we should do in any case, since it may be that despite our running a significant risk in accepting it, C is in fact true. In that case there may be no significant epistemic risk in accepting P since P can then be a deliverance of reliably functioning perceptual faculties in an external material world that is receptive to them and much as we take it to be, with the result, so we may take it, that P is thereby known.

What, Wright (2004) conceded, *is* at risk as a consequence of the riskiness of accepting C is rather the higher-order claim that (4) itself is known, more specifically, that we can *knowledgeably lay claim* to know P. Rather, the position is, in effect, that while in general the overall conditions sufficient to enable the acquisition of perceptual knowledge may obtain, we have only an entitlement to suppose so. That is not inconsistent with (4) — with our actually possessing such knowledge of P. But it is inconsistent with our knowing or having a balance of evidence, as opposed to being entitled to suppose, that we do. All the same, with that conceded, we can still be rationally entitled to claim that that is indeed our fortunate situation.

In summary: Wright (2004) conceded that there is indeed ‘leaching’ of entitlement upwards from cornerstones but contended that what is thereby displaced is not — better: may not justifiably be affirmed to be — our knowledge of, or perceptual justification for accepting P, but the *second-order* epistemic credentials of such propositions, which ought consequently to be re-characterised as matters of entitlement.

And with that response Wright (2004) was content to leave the matter.⁵

3. Varieties of (Epistemic) Risk

We need a caveat before proceeding to review the stability of this response. The term ‘risk’, as it occurs in contexts of advice, admonition, or prudential reasoning is customarily taken to connote *the probability of an unwanted event or situation*, where this probability is estimated

⁴ In an Arché seminar discussion in the early 2000s.

⁵ For critical discussion of Wright’s response independent of the argument to follow, see McGlynn (2017).

on a relevant body of evidence. Thus, the degree of risk associated with a proposition Q describing an unwanted event or situation is customarily thought of simply as Q 's degree of *evidential* probability (cf. Hansson 2018 and Ebert *et al.* 2019). This usage is, for instance, exemplified by the statement that, on the current medical evidence, the risk that a smoker's life is shortened by some smoking-related disease is about $\frac{1}{2}$ (cf. Hansson 2018).⁶ Sometimes people speak of the risk that one runs in accepting a proposition, P , as true, or acting on the assumption that P is true. In this case the unwanted situation is one in which P is actually false. So, the risk that one runs in accepting or acting on the assumption of P is actually the risk associated with P 's being false.

However, it is by no means self-evident that our intuitive judgements about risk in general track — or ought to track — probability. The recent literature⁷ has supported the credentials of two interestingly non-probabilistic notions of risk. Duncan Pritchard has argued that in a wide class of cases, the risk one runs in accepting, or acting on the assumption of a certain proposition, P , is best assessed not in terms of P 's probability — whereby the smaller the probability, the higher the risk of acceptance — but in modal terms: roughly by how *safe* the assumption of P is — that is, by how similar to the actual world are the nearest worlds in which P fails: the more similar these worlds to actuality, the higher the risk of accepting P . Thus, on this view, the risk of accepting P is higher than the risk of accepting Q just in case the worlds most similar to actuality in which P is untrue are more similar to actuality than the worlds most similar to actuality in which Q is untrue.

Pritchard attempts to motivate his proposal by appeal to a scenario in which our intuitions of degree of risk are supposed to align inversely with the extent of the departure it involves from actuality rather than with chance. In our view there are serious problems with the details of the example.⁸ However we are sympathetic with the general suggestion that our estimate of the degree of risk involved in an assumption often has more to do with the perceived extravagance of the nearest possible scenarios in which it fails than with any estimate, however rough, we feel able to make of the numerical probability of its obtaining.

⁶ A refinement of this conception, adopted in professional risk management and economics, interprets 'risk' as referring to the *expectation value* of a possible negative event, which is the product of the event's probability and some measure of its severity (cf. Hansson 2018). To avoid sterile complications, we set aside this more technical notion.

⁷ See Pritchard (2015, 2016) and Ebert *et al.* (2020).

⁸ Here is (an edited version of) Pritchard's (2015 and 2016) principal example:

Bomb 1: An evil scientist has rigged up a bomb hidden in a populated area. There is no way of discovering the bomb. It will detonate only if a certain set of numbers comes up on the next national lottery draw. The odds of these numbers appearing is fourteen million to one. It is impossible to interfere with this lottery draw.

Bomb 2: An evil scientist has rigged up a bomb hidden in a populated area. There is no way of discovering the bomb. It will detonate only if three highly unlikely events all obtain. First, the weakest horse in the field at the Grand National, Lucky Loser, wins the race by at least ten furlongs. Second, the worst team remaining in the FA Cup draw, Accrington Stanley, beats the best team remaining, Manchester United, by at least ten goals. Third, the Queen of England spontaneously chooses to speak a complete sentence of Polish during her next public speech. The odds of this chain of events occurring are fourteen million to one. It is impossible to interfere with the outcomes of any event in this chain.

Let P state that the bomb will detonate. Pritchard asserts that although P is equally likely in both cases, the assumption that P is false — that the bomb will not detonate — is intuitively appreciably riskier in Bomb 1 than in Bomb 2. Pritchard explains this intuition by suggesting that P appears safer in Bomb 1 than Bomb 2. But it is unclear to us that this explanation is correct. It is a mere stipulation by Pritchard that the event that would trigger the detonation in Bomb 1 can actually meaningfully be said to have the same degree of chance as the sequence of events that would trigger the detonation in Bomb 2. It is consequently quite unwarranted to assume that any sense of diminished riskiness of P 's being true that folk do indeed have in contemplating Bomb 2 is sensitive to this merely stipulated, possibly meaningless, equality.

Pritchard's modal (safety) proposal contrasts with a third — the so-called “Normic” account of risk, canvassed by Ebert, Smith and Durbach (2020). According to this account, the risk of a proposition *P* is determined by the degree of *normality* of the most normal worlds in which *P* is true. The more normal these worlds are — that is to say, the closer they are to the maximally normal worlds — the higher the risk of *P*. On this view, the risk of *P* is higher than the risk of *Q* just in case the most normal worlds in which *P* is true are more normal than the most normal worlds in which *Q* is true.

Three points should be noticed: first, on the normic account, the maximally normal worlds need not include the *actual* world. As Ebert, Smith and Durbach stress, “truth doesn't entail normalcy and there is nothing contradictory in saying ‘*P* is true but *P* wouldn't normally be true.’”⁹ Second, the normality of a world *w* is understood as determined by the extent to which the respects in which *w* differed from the maximally normal worlds would call for special explanation, were *w* to obtain. (Thus for example, while the probability of Martin's ticket winning the EuroMillions Jackpot lottery might be vanishingly small, it would call for no special explanation — no significant departure from normality — were it to do so. Somebody has to win).¹⁰ Third, the ordering of worlds by degree of normality is to be understood as restricted to worlds in which our actual background evidence holds, and should exclude worlds which are inconsistent with this background evidence.

Ebert, Smith and Durbach offer no intuitive example in which the normic account diverges from the other two accounts and seems intuitively superior. But we will forbear from attempting to add here to the debates about these conflicting proposals. Indeed, there is a recent tendency to some agreement in the literature that our pre-theoretic judgements about risk do not provide a basis for identifying any one thing as deserving to be regarded as *the* correct notion of risk.¹¹ Rather, our understanding of risk seems to be ‘open textured’ to a degree, admitting of more than one explication or precisification, so that a form of pluralism about risk may be the right way to think about the notion.

What is striking for the immediate purpose, however, is that neither the modal nor the normic accounts of risk promise to be of any service to an attempt at consolidation of the Leaching worry. The point is especially clear on the modal account. For suppose that there actually is indeed an external material world, that is very much as we judge it to be on the basis of ordinary perception. Then the nearest worlds in which there is no such material reality are massively dissimilar to the actual world. Accordingly, by the modal account, we run no significant risk in presuming the relevant cornerstone proposition to be true. Hence to presume, with the sceptic, that there is justification for thinking there is such a significant risk is, on the modal account, implicitly to suppose there is justification for an assumption

⁹ Ebert, Smith and Durbach (2020), p. 444.

¹⁰ The reader may well feel that this is not a great result for the normic view, since — worlds in which Martin's ticket wins being roughly no less normal than worlds in which it doesn't, — it seems we'd have to regard the alternatives, Win or Lose, as equally risky on the normic account. Of course it is odd to describe a good outcome as a ‘risk’ at all; but we could easily concoct an example where winning is a bad result. As Matt Jope has reminded us, a somewhat macabre but vivid example where normic considerations arguably provide a better recovery of our intuitions about risk than probabilistic ones can be provided by an embellishment of Shirley Jackson's short story *The Lottery* - (The New Yorker, June 26 1948) - wherein a rural community runs a large lottery each year in the belief that full participation alone will serve to ward off a poor harvest and consequent famine in the following year. The sting is that in order to appease the agrarian Gods, the ‘winner’ must be put to death. The sense of risk had by the participants in such a setting would, plausibly, be much stronger than can be explained by the numerical probabilities alone. So much would be predicted by the normic account since for each participant, however statistically improbable it may be, ‘winning’ would be a wholly normal occurrence — someone has to ‘win’ — and thus would constitute a significant risk.

¹¹ Cf. Bricker (2018), Moretti (2019) and Ebert, Smith and Durbach (2020).

inconsistent with the truth of the cornerstone proposition—and hence to suppose that we are justified in *denying* that there is an external material world. Which of course we are not.

The situation on the normic account is a little more complex. On the normic account, in order for an acceptance of the material world cornerstone to emerge as significantly risky, it has to be the case that there are worlds that include our actual ordinary background evidence, which are pretty normal and in which the cornerstone fails. Ask first: what would that require of “ordinary background evidence”? If it includes, as externalists will want to say, perceptual knowledge of our immediate material environment, then obviously there won’t be any normal worlds including *that* background evidence in which the cornerstone fails. So in order for our “ordinary background evidence” to be understood as consistent with failure of the cornerstone, an argument for thinking that acceptance of the cornerstone is risky on the normic conception will have to work with an internalist, “highest common factor” conception of “ordinary background evidence”.

Well and good, it may be thought. Isn’t that in any case the conception of evidence to which one who, like Wright, concedes that the sceptical paradoxes have shown that the cornerstones are beyond evidence, is committed? To be sure. But now there are problems in explaining how acceptance of the cornerstones is risky by the lights of the normic conception. In fact there are two separate problems. The first is that our conception of what counts as abnormality, and thus as calling for special explanation, is constrained by a conception of the natural order of things which *presupposes* the usual cornerstones. Events and states of affairs are determined to be normal or abnormal relative to what is expectable, or anyway unremarkable, within explanatory frameworks for which the cornerstones serve as just that, and the explanations which abnormal occurrences call for are accordingly explanations in which the cornerstones are taken for granted. It follows that both the cornerstones and their sceptical negations are beyond the scope of kind of explanation which, on the normic account, abnormality constitutively demands—the cornerstones because they are presupposed in the relevant explanatory frameworks, and their sceptical negations because they are inconsistent with those frameworks. The normic conception of risk is thus essentially inapplicable to both.

But there is a second, more intuitive point. Again, in order for it to rank as risky to accept a cornerstone *C* on the normic conception of risk, we have to make sense of the idea that *C* can fail in relatively normal worlds in which we have our actual (internal) evidence. And that requires that were *C* to fail, then that would call for no special explanation. But that seems intuitively preposterous. Surely, if the course of our actual experience were to have been what it has been and were to continue as, broadly, it always has and yet there were no external material reality, that would demand a Whale of an Explanation! If God speaks to you in a dream and advises you that such is the situation, saying nothing further, your reaction will not be, “Oh well. That was always on the cards” in the spirit of “Someone has to win”.

We may conclude that if there is a development of the Leaching worry that does successfully exploit a notion of risk so as to do significant damage to the Wright (2004) entitlement proposal, it must work with a probabilistic conception of risk. Which brings us to the principal focus of this paper.

§4 The Leaching worry revived?

Moretti (2020) gives a detailed argument that, notwithstanding Wright’s original response, the Leaching worry can indeed be vindicated by orthodox probabilistic reasoning. In particular he argues that, once the pre-theoretical notion of significant epistemic risk is construed probabilistically and the Entitlement theorist allows that we are at significant epistemic risk in accepting cornerstones, significant epistemic risk must be acknowledged to transfer from acceptance of the cornerstones back to acceptance of the various relevant kinds of quotidian propositions about the external world, others’ mental states, and the past that we

take ourselves habitually and continually to come to know — and, to stress, that the degree of risk transferred is sufficient to undermine both that knowledge and the idea that we are evidentially justified in accepting such quotidian propositions

Here is the argument. Let Pr be a probability function interpreted subjectively. $Pr(Q)$ is to express one's degree of rational confidence that Q . $Pr(Q|R)$, likewise, is to express one's degree of rational confidence that Q conditional on one's learning R . Consider again an empirically testable proposition P — say that there is a hand in front of your face — and a correlated sceptical alternative SH — say that there is no external material world but that a Mischievous Demon, who controls all your experiences and persistently maintains the illusion of an external world, is causing it to appear to you that there is a hand in front of your face. Note that we have chosen SH to be incompatible with P but specific enough to entail that you have a deceptive experience as of P 's being the case.¹² Let E be the proposition that you have an experience as of P 's being true.

So: we have that P entails $\neg SH$, and SH entails E . Roger White (2006) has shown that provided those conditions are met, and if (1) and (2) following are accepted:

(1) $Pr(SH) > 0$, (which seems entirely reasonable, since the sceptical scenario is surely not *certainly* false);

(2) $Pr(E) < 1$, (which again seems entirely reasonable as a prior probability, since E describes a future contingency),

then the probability that P is true, given that the experience depicted in E does indeed duly occur, is *strictly less than* the probability that the sceptical scenario depicted by SH does not obtain; that is, that

$$(*) \ Pr(P|E) < Pr(\neg SH)^{13}$$

(which again, on reflection, seems entirely reasonable since among possible scenarios in which the E -predicting sceptical hypothesis is false, — because for instance there is indeed

¹² We emphasise to the reader that this feature is essential to the argument. The reasoning to follow requires the assumption that we are working with a sceptical hypothesis, SH , that predicts your actual evidence for P , or at least a fully convincing simulacrum of it. Suppose, for example, that P is 'There is a hand in front of my face', asserted on the basis of current perceptual experience, that E reports such an experience and that $\neg SH$ is simply 'There is a material world'. In this case P entails $\neg SH$, but SH — i.e. 'There is no material world' — doesn't entail E — that is, that I undergo an experience as if there is a hand in front of my face. In such a case the following proof doesn't apply, and the revamped Leaching worry about to be developed will be pre-empted.

¹³ Here is White's proof:—

Since $Pr(SH) > 0$ and SH entails E , $Pr(E|SH) = 1$. From this, since $Pr(E) < 1$, $Pr(E|SH) > Pr(E)$. Thus, $Pr(SH|E) > Pr(SH)$. Thus,

(i) $Pr(\neg SH|E) < Pr(\neg SH)$.

Since $Pr(SH) > 0$ and SH entails E , then $Pr(E) > 0$. From this, given that P entails $\neg SH$, it follows that

(ii) $Pr(P|E) \leq Pr(\neg SH|E)$.

From (i) and (ii), by transitivity, $Pr(P|E) < Pr(\neg SH)$. QED.

White originally presented this nice observation as a problem for Jim Pryor's (2000) *perceptual dogmatism* — for the idea that the evidential force of ordinary perceptual experience is, as a default, presuppositionless. There is an irony, therefore, in the fact that the very same theorem may be deployed against the attempt, via the notion of Entitlement, to shore up the opposed perceptual *conservatism*, whose signature claim is to grant the sceptic the presuppositional role of cornerstones, against the sceptical paradoxes (as in Wright 2004 and 2014).

an external material world — there will be instances where P is false even though the experience depicted by E occurs.)

Proposition (*) says that if S has an experience of the kind depicted by E , S 's degree of confidence on that basis in P , if rational, must remain strictly smaller than S 's *prior* confidence in $\neg SH$. Thus, S 's learning that E is true can make S rationally confident of P 's truth only if S is antecedently *more* confident of $\neg SH$'s truth, i.e., that the sceptical scenario is false.

Well then, let's now suppose specifically that Pr is a function purely of *evidential* probability.¹⁴ The degree of epistemic risk involved in the assumption that you are in the 'good case' — i.e. that SH is false — can be identified with $Pr(SH)$. And the degree of epistemic risk that P is false, given E , can be identified with $Pr(\neg P|E)$. It is straightforward to show that White's result above

$$(*) \quad Pr(P|E) < Pr(\neg SH)$$

is equivalent to:

$$(Quasi-leaching) \quad Pr(\neg P|E) > Pr(SH).^{15}$$

(*Quasi-leaching*)¹⁶ states that, on assumptions (1) and (2) above, the probability that P is false given E — so the degree of evidential risk one runs in accepting P , given E — is strictly higher than the probability that the sceptical hypothesis is true. So you undertake a greater risk in accepting P on the basis of the experience depicted by E than you undertake in rejecting the sceptical hypothesis. You are more likely to be mistaken in thinking you have a hand on the basis of your current visual and kinaesthetic experience as of a hand than you are in discounting the Mischievous Demon scenario!

More explicitly: let r be an *evidential threshold* of significant epistemic risk, that is, for any proposition Q and evidence E , if $Pr(Q|E) > r$, the risk that Q is false for a subject, S , whose only germane evidence is E is so high that S 's accepting that proposition would be evidentially unjustified. Then (*Quasi-leaching*) entails that if $Pr(SH) > r$, then $Pr(\neg P|E) > r$. That is to say, if S 's accepting $\neg SH$ on the basis of a non-evidential entitlement — and so prior to learning E — is *evidentially* unjustifiably risky, S 's accepting P upon acquisition of the evidence E is also evidentially unjustifiably risky.

It is a matter for judgement whether this result directly contradicts Wright's response to the original leaching worry or whether it merely drains away its purported significance. Wright's response was that while your *evidential* deficiency in relation to the material world cornerstone should be acknowledged, you still retain a rational entitlement to accept it, and within the epistemic framework provided by that acceptance, are therefore entitled to regard your experience as powerfully evidential, indeed as knowledge-conferring, for P . If there is a risk in accepting that SH is false, that risk, the suggestion was, properly transmits not to the

¹⁴ Note that evidential probability need not be represented as a *conditional* probability. The relevant evidence can be thought of as *implicit* in the unconditional probability function.

¹⁵ This is so because $Pr(P|E) = 1 - Pr(\neg P|E)$ and $Pr(\neg SH) = 1 - Pr(SH)$.

¹⁶ We are speaking of "Quasi-leaching" because although this thesis arguably does not directly imply that the risk of accepting P on the basis of E is an *inheritance* of the risk of accepting that the sceptical hypothesis is false (which was part of the original metaphor of "leaching upwards"), the apparent awkwardness of the result for the Entitlement theorist — the challenge to the rational coherence of the more favourable epistemic attitude to P than to SH to which the theorist aspires — is not diminished on that account.

acceptance of P itself but only to the second-order claim that one knows P on the basis of one's experience — in effect, to the claim that one's experience of E is a source of knowledge. Moreover that risky claim can still be one to which one is rationally entitled. However the sting in (*Quasi-leaching*) is that it is, seemingly, *the acceptance of P itself* on the basis of E and not merely the second-order claim to know it, that is disclosed as evidentially dangerous. Knowing P is not consistent with undertaking a significant epistemic risk in accepting P . So even if one enjoys nevertheless a rational entitlement to claim to know P on the basis of E , the result established in (*Quasi-leaching*) has not gone away that to *act on* that alleged entitlement — that is, actually to accept P — is yet *more* epistemically risky than to accept the cornerstone, the riskiness of which was acknowledged by the Entitlement proposal from the start. That proposal, it appears, must therefore contrive some way of blocking White's result.

§5 Three responses

We shall canvass three possible responses on behalf of the Entitlement theorist, the first of which we will argue is beside the point, while the second and third may offer better prospects for stabilising Wright's proposal.

(i) First, it might be suggested that since there is absolutely no evidence in support of the kind of sceptical hypotheses we are concerned with, their evidential probability ought actually to be reckoned to be very low, so that our trust in their negations should not be considered to be epistemically risky in the first place. In that case, the leaching worry is misconceived from the get-go: there is no significant epistemic risk that we run by accepting cornerstones and that might leach upwards to our routine acceptance of quotidian propositions that entail them.

Indeed, it might seem that there is a standpoint from which this response is fully justified. Reflect that if $\Pr(SH)$ were very low, then since $\Pr(SH) + \Pr(\neg SH) = 1$, $\Pr(\neg SH)$ should be *extremely high*. It follows that one who offers this first response is committed to holding that we have an overwhelming justification for accepting the negation of SH — an overwhelming justification for holding that there is an external material world, that other minds are real, etc. How might that claim be backed up? Well, isn't that exactly the position of the kind of Moorean *liberalism* defended by Pryor and others?^{17,18}

¹⁷ Pryor (2000).

¹⁸ It is, by the way, worth noting that the Leaching worry as developed in the preceding section is unmitigated if we consider the epistemic predicament in which we are left by cogent sceptical argument to be one where we lack evidence both for a cornerstone and a sceptical hypothesis inconsistent with it, so that our epistemic predicament appears to be such that both $\Pr(SH)$ and $\Pr(\neg SH)$ are largely *undetermined* by our evidence. We can straightforwardly refashion the notion of significant epistemic risk to include situations of *lack* of evidence. Let's stipulate that S 's accepting Q is significantly epistemically risky if the probability of $\neg Q$ is *undetermined* by S 's evidence within a wide interval of values the least upper bound of which exceeds r (where r , as above, is a threshold that, if exceeded, prevents S from justifiably believing Q). On this interpretation of significant epistemic risk, saying that S 's accepting $\neg SH$ is epistemically risky is saying, not that $\Pr(SH)$ exceeds r but, rather, that $\Pr(SH)$ is indeterminate within a wide interval whose least upper bound exceeds r . Since S lacks evidence in favour or against $\neg SH$, it looks very plausible that S 's accepting $\neg SH$ is epistemically risky in this redefined sense. Moreover, it is easy to see that the epistemic risk of $\neg SH$ continues to infect the acceptance of P even when it is characterised in this way. Since $\Pr(SH)$ is indeterminate within an interval whose least upper bound exceeds r , *Quasi-leaching* entails that $\Pr(\neg P|E)$ must also be indeterminate within an interval whose least upper bound exceeds r . As before, once S 's accepting $\neg SH$ on the basis of a non-evidential entitlement is allowed to be epistemically risky, S 's accepting P upon learning E must also be no less — indeed more — epistemically risky.

Such is indeed the Moorean liberal position, but invoking it here is doubly irrelevant. First, our concern is with the stability of the kind of response to scepticism proposed in the opposed *conservative* spirit that already concedes that the sceptical arguments do indeed make a case that the cornerstones are beyond evidence. But second, there is anyway an implicit misunderstanding of White's result. That result is that the probability one attributes to P , given E , is constrained to be lower than the *prior* probability one assigns to the negation of SH . Thus even if, as Moorean liberalism supposes, the cornerstones have accumulated overwhelming evidential support from centuries of humankind's quotidian experience, that says nothing about the probability they should be reckoned to have in advance of any particular experience. But it is that probability which an entitlement theorist who grants that the acceptance of cornerstones is irretrievably significantly risky, thereby arguably concedes to be significantly below 1, thereby conceding the first premise of White's reasoning,

(1) $Pr(SH) > 0$.

(ii) The second response to the revamped Leaching worry pursues one way in which White's first premise may be challenged. (The third response will consider another.) It runs as follows. Rational entitlement is entitlement to "take for granted", to trust "implicitly", as we say; it is to bracket any doubt. But if determining risk is to be calculation of probabilities, and cornerstones are to be included within the scope of the determination, then within a Bayesian framework we need to assign priors to cornerstone propositions. True, when Pr is a function of *evidential* probability, White's reasoning assumed only that such propositions are not evidentially certain. But is it satisfactory to construe $Pr(C)$ for these purposes purely as a function of evidential probability? It is suggestive that Wittgenstein's last notes, from which Entitlement theory takes inspiration, are entitled *On Certainty*, and it was not evidential certainty that Wittgenstein had in mind! The second response to the revamped Leaching worry will contend that when S is rationally entitled to trust $\neg SH$, this trust should be unreserved: S 's entitlement is to regard $\neg SH$ as assured and to treat it accordingly as *beyond doubt*. Consequently, we should set $Pr(\neg SH)$ as equal to 1 and $Pr(SH)$ as equal to 0, thus cutting the ground from under the proof of (*Quasi-leaching*). Under that assignment, an acceptance of $\neg SH$ should no longer be regarded as epistemically risky and there will accordingly be no significant epistemic risk that might be quasi-leached to P . When our entitlements are conceived as *legitimate certainties*, nothing prevents a subject from justifiably believing P on the basis of E , indeed from claiming knowledge on that basis, with a corresponding certainty.

This is liable to provoke the immediate rejoinder that to respond in this manner is simply to forget the way in which Pr needs to be interpreted in order to model epistemic risk in the first place. The kind of risk we are concerned with — the kind of risk that, after all, Wright's original response to the leaching worry does implicitly concede that we do indeed run in accepting sceptically-challenged cornerstones—is essentially that of *going beyond the evidence*, and in order to be at the service of modelling *this* notion of epistemic risk, Pr accordingly needs to be a function purely of evidential probability. So, Pr can account only and specifically for S 's degrees of rational confidence based on the evidence available to her. The alleged fact that S is *non-evidentially* entitled to accept $\neg SH$ with complete assurance cannot be allowed to imply anything about the value of $Pr(\neg SH)$ properly so interpreted, still less that it approximates 1.

The rejoinder is fair against the letter of Wright (2004) *ad hominem*. But a proponent of the second response should reply that, once entitlement is accepted as a form of genuinely

epistemic — though non-evidential — warrant, it is no longer satisfactory to interpret the notion of significant epistemic risk purely as evidential risk. Rather the epistemic risk one runs in accepting *P* should be computed as a function of one's total epistemic situation, encompassing both one's evidential and non-evidential warrants. It can still be acknowledged that one runs an evidential risk in accepting cornerstones. But there is now no *normative* punch associated with that acknowledgement. The acceptance of things which one is rationally warranted in treating as certain, if that is indeed the pay-off of entitlement, cannot coherently be conceived as the running of risks in any sense that implies epistemic culpability.

(iii) The foregoing is one way of implementing the more general idea that the cornerstones and other entitlements should somehow occupy a special status when it comes to the probabilistic computation of risk. But a sneaking sense may persist that to pursue the general idea in that way — to treat the cornerstones as certain and sure — is somehow to disrespect scepticism, to undervalue the intellectual achievement of the best sceptical paradoxes. Against that, it may be said that what is being mooted is an entitlement to *treat* the cornerstones as certain, not to *feel* that they are certain — the certainty that Hume, at least in the study, found it so difficult to feel. However, we will not pursue this issue here. A third way to respond to the revamped leaching worry avoids investing in the putative *certainty* of entitlements in either sense. It starts with acknowledging that there is indeed an open question about what exactly our epistemic entitlements are best conceived a licensing us to do — what attitude or action is rationalised thereby. We have so far been assuming, with Wright (2004 and elsewhere) that entitlement should be conceived as something in the ballpark of a license to *accept* or *trust* in the hinge proposition “There is a material world”. But, on reflection, it is not obvious that this kind of interpretation is indeed even in the best ballpark. Advocates of entitlement theory can contend that in this and other cases of this type, one's entitlement is best construed in a way inspired by David Lewis' classic contextualist manifesto¹⁹ as entitlement to *ignore* $\neg SH$ — or more exactly, to ignore both *SH* and, consequently, $\neg SH$.

This idea should not, we suggest, seem terribly odd or unnatural. An entitlement of this type, to disregard both a sceptical hypothesis and the cornerstones that it attacks, would simply sanction what a normal subject would probably already do. In daily life and science, people ignore sceptical possibilities even when they are aware of them. An epistemologist, fresh from giving a talk on Descartes and attempting to remember on what level she earlier parked her car in the airport parking structure, may experience mounting anxiety but one thing she won't be anxious about will be the possibility that her memory might have been tricked by a Cartesian demon, or indeed that she has no car nor any other material possessions for that matter, including her physical body. Similarly, the physicists who have recently detected gravitational waves have certainly ignored the conjecture that the Matrix might have fabricated their data, though some of them at least are very likely have watched the film “The Matrix”, or are anyway familiar with the story told in it.

Someone might suggest that in cases like these, although admittedly at some level aware of them, the subjects are able to ignore sceptical possibilities only because they don't actually *attend* to the scenarios involved, which remain, like a prospectively uncomfortable dental appointment still some weeks away, at the periphery of their thoughts; and that if the imagined epistemologist or the physicists were actually to focus on the relevant sceptical hypotheses, then of course they could no longer possibly ignore them. If that were correct, the proposal that we have entitlements to ignore sceptical scenarios would be pointless, since we wouldn't be able to cash the license such entitlements provided, once we confronted any

¹⁹ Lewis (1996).

specific sceptical challenge. In order to be of any counter-sceptical use, our epistemic entitlements have to be things we can exploit precisely when we *attend* to sceptical conjectures. Rights ain't worth a lot if we cannot appeal to them in exactly the situations when they might benefit us.

An emblematic example of the idea at work in that train of thought is Lewis's Rule of Attention, according to which "a possibility that is not ignored at all is *ipso facto* not properly ignored"²⁰ Lewis thought — or seems to have thought — that this principle should strike one as self-evident. However, for the kind of ignoring that is here germane—ignoring not as ignorance but as a kind of active disregard — the Rule of Attention is not merely not self-evident but manifestly tendentious. The relevant point is succinctly made by Michael Williams:

If (rudely or quite understandably) I ignore you at a party, this is not because I don't realise you are there. On the contrary, I have to know you are there to ignore you. So too in epistemic matters. I can (properly or improperly) ignore — i.e. not take into account — possibilities of which I am fully aware.²¹

Since we are in this sense psychologically capable of ignoring things we attend to, one may wonder — *pace* Lewis — in what cases we can *properly* ignore things that we attend to — in particular, sceptical possibilities like *SH*. The answer being suggested is that we can properly ignore sceptical scenarios — in the sense of legitimately refusing to take them into account in our epistemic evaluations — when we are *epistemically entitled* to do so.

How does entitlement, so conceived as to license deliberate disregard, provide the resources to defuse the revamped Leaching worry? Very directly. White's result and its corollary, (*Quasi-Leaching*), depend as noted on our taking *SH* and its kin to lie within the range of propositions over which *Pr* is defined. There is no reason, once we are entitled to ignore certain possibilities, why that range should be taken to include such possibilities — why *SH* should be included in the range of propositions over which a fully rational subject, *S*, distributes her degrees of confidence. Clearly, if *S* is entitled to ignore the possibility *SH*, *S* is entitled not to include *SH* (and thus $\neg SH$) in the relevant probabilistic algebra. And if *SH* isn't included, "*Pr*(*SH*)" no longer has a denotation, and the first assumption of White's proof, that

$$(1) \Pr(SH) > 0$$

is once again undercut.²²

References

- Blome-Tillmann, M. 2014 *Knowledge and Presuppositions*. Oxford: Oxford University Press.
 Blome-Tillmann, M. 2009. 'Knowledge and Presuppositions'. *Mind* 118: 241-294.
 Bricker, A. M. 2018. 'Do judgements about risk track modal ordering?' *Thought* 7: 200-208.
 Dretske, F. 2005. 'The Case against Closure'. In Steup, M. and Sosa, E. (eds.), 2005: 27-40
 Ebert, P. A, Smith, M. and J. Durbach 2020. 'Varieties of Risk.' *Philosophy and Phenomenological Research* 101: 432-455.
 Hansson, S. O. 2018. 'Risk'. In E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*.
<https://plato.stanford.edu/archives/fall2018/entries/risk/>

²⁰ Lewis (1996), p. 559.

²¹ Williams (2001) at p.16. See also Blome-Tillmann (2009 and 2014).

²² We are very grateful to Philip Ebert, Matt Jope and Martin Smith for comments on early drafts of this paper, and for seminar discussion of the penultimate version by the members of the AHRC-funded *Varieties of Risk* project co-hosted by Stirling and Edinburgh from September 2020 to March 2023.

- Hawthorne, J. 2005. 'The Case for Closure'. In Steup, M. and Sosa, E. (eds.), 2005: 40-56
- Lewis, D. 1996. Elusive knowledge. *Australasian Journal of Philosophy* 74: 549-567.
- McGlynn, A. 2017. 'Epistemic Entitlement and the Leaching Problem'. *Episteme*, Vol. 14 (1), March 2017, 89-102 DOI: <https://doi.org/10.1017/epi.2015.63>
- Moretti, L. 2020. 'Entitlement, Epistemic Risk, and Scepticism'. *Episteme*. First View, 1 – 11. DOI: <https://doi.org/10.1017/epi.2019.27>
- Pritchard, D. 2015. 'Risk'. *Metaphilosophy* 46 (3):436-461
- Pryor, J. (2000). 'The Skeptic and the Dogmatist'. *Noûs*, 34(4), 517-549.
- 2004. 'What's wrong with Moore's argument?' *Philosophical Issues*, 14, *Epistemology*: 349-378.
- Steup, M. and Sosa, E. (eds.), 2005, *Contemporary Debates in Epistemology*, Malden, MA: Blackwell.
- White, R. 2006. 'Problems for Dogmatism'. *Philosophical Studies* 131: 525-257.
- Williams, M. 2001. 'Contextualism, Externalism and Epistemic Standards'. *Philosophical Studies* 103: 1–23.
- Wright, C. 2004. 'On Epistemic Entitlement: Warrant for Nothing (and Foundations for Free?)'. *Proceedings of the Aristotelian Society, Supplementary vol.* 78: 167–212.
- 2007. 'The Perils of Dogmatism' in *Themes from G. E. Moore: New Essays in Epistemology and Ethics*, edited by Susana Nuccetelli and Gary Seay, Oxford University Press: 25-48.
- 2012. 'Warrant Transmission and Entitlement', Part IV of "Replies" in Annalisa Coliva (ed.) *Mind, Meaning and Knowledge: Themes from the Philosophy of Crispin Wright*, Oxford University Press: 451-486.
- 2014. 'On Epistemic Entitlement (II): Welfare State Epistemology'. In E. Zardini and D. Dodd (eds.), *Scepticism and Perceptual Justification*, Oxford University Press: 213-247.