

ON EPISTEMIC ENTITLEMENT

By Crispin Wright and Martin Davies

I—Crispin Wright

WARRANT FOR NOTHING
(AND FOUNDATIONS FOR FREE)?

My *life* consists in my being content to accept many things
(Wittgenstein *On Certainty* §344)

ABSTRACT Two kinds of epistemological sceptical paradox are reviewed and a shared assumption, that warrant to accept a proposition has to be the same thing as having evidence for its truth, is noted. ‘Entitlement’, as used here, denotes a kind of rational warrant that counter-exemplifies that identification.¹ The paper pursues the thought that there are various kinds of entitlement and explores the possibility that the sceptical paradoxes might receive a uniform solution if entitlement can be made to reach sufficiently far. Three kinds of entitlement are characterised and given *prima facie* support, and a fourth is canvassed. Certain foreseeable limitations of the suggested anti-sceptical strategy are noted. The discussion is grounded, overall, in a conception of the sceptical paradoxes not as directly challenging our having any warrant for large classes of our beliefs but as crises of intellectual conscience for one who wants to claim that we do.²

I

Two Kinds of Sceptical Paradox. Call a proposition a *cornerstone* for a given region of thought just in case it

1. The term is already in use in contemporary epistemology in a number of contrastive senses. My use of it contrasts in particular—though it also has points of contact—with that of Tyler Burge in a number of important recent articles (see e.g. Burge [1993]). Such overlap in terminology is unfortunate but, given that English has only so many expressions for norms of doxastic acceptance, all of which are already in use with multiple connotations, it is too late to hope to avoid it.

2. The paper originates in ideas that go back to my Henriette Hertz British Academy lecture (Wright [1985]) which shared the root idea that an attractive response to scepticism might draw on the possibility of non-evidential warrant. The major strategic contrast with the present proposals is in how such warrant is conceived as possible. In the lecture, I proposed that at least some ‘cornerstones’ might be regarded as defective in factual content and that acceptance of them might accordingly be freed from the requirements of evidence that I took to be characteristic of the factual. In the present discussion, non-factuality is no longer assigned a role in making a case that rational acceptance need not be evidence-based.

would follow from a lack of warrant for it that one could not rationally claim warrant for *any* belief in the region. The best—most challenging, most interesting—sceptical paradoxes work in two steps: by (i) making a case that a certain proposition (or restricted type of proposition) that we characteristically accept is indeed such a cornerstone for a much wider class of beliefs, and then (ii) arguing that we have no warrant for it.

The ‘best’ such paradoxes are, I think, of essentially two kinds, though they each allow of minor variations of detail. The first—what we may call *Cartesian*—makes a case that it is a cornerstone for a large class of our beliefs that we are not cognitively disabled or detached from reality in a certain way—the scenarios of a persistent coherent dream or hallucination, persistent deception by a *malin génie*, the envatment of one’s disembodied brain, and ‘The Matrix’ are examples of such detachment—and then argues that we have no warrant to discount the scenario in question. So the upshot is a challenge to our possession of warrant for any of the large class of dependent beliefs in question.

There are various ways a Cartesian sceptical argument may support its two ingredient lemmas. The details cannot concern us now—otherwise we won’t get to the issues I want to get to. But we do need to register a point about how the second lemma—that we have no warrant to discount the relevant scenario—is usually supported. Suppose I do have warrant to discount the suggestion that I am right now in the midst of a sustained and coherent dream? Well, if I have such a warrant, how did I get it? The proposition that I am not right now suffering such a dream is, broadly speaking, an empirical one, so any warrant I have for it must presumably consist in empirical evidence, acquired by executing some appropriate empirical procedure. However—the sceptical argument says—evidence acquired as the result of an empirical procedure cannot rationally be regarded as any stronger than one’s independent grounds for supposing that the procedure in question has been executed properly. For instance, measurement-based evidence to regard the edge of my desk as near enough 1.75 metres long cannot rationally be regarded as any stronger than my independent warrant to suppose that the measuring procedure was carried out to appropriate tolerances, using a properly calibrated tape measure, and the results carefully observed, etc.

A fortiori, then, evidence for the proposition that I am not now dreaming, acquired as a result of executing some appropriate empirical procedure, cannot rationally be regarded as any stronger than my independent warrant for thinking that the relevant procedure was properly executed, and hence for thinking that it was executed *at all*—ergo: that I did not merely dream its execution! So it appears that my acquiring a warrant by empirical means for the proposition that I am not now dreaming requires that I *already have* a warrant for that same proposition. So I cannot ever acquire such a warrant (for the first time.)

No doubt that reasoning is very discussible. The only point I want to call to your attention at this stage is that—strictly—its conclusion falls short of the needed sceptical lemma. The conclusion is that I can't get evidence to discount the supposition that I'm right now in the midst of a sustained and coherent dream. The needed lemma is that I *don't have* any warrant to discount that supposition. The lacuna will matter if it's possible to have warrant (for an empirical proposition) which does not consist in the acquisition of empirical evidence. I conjecture that this lacuna will be left unfilled by all likely ways of arguing that a Cartesian sceptical scenario cannot warrantably be discounted.

The second principal genre of (interesting) sceptical paradox is typified by Hume's inductive scepticism. Here there is no play with a scenario of cognitive disablement or dislocation. Rather, the sceptical argument makes a case that our epistemic procedures involve a vicious circle. The challenge posed by inductive scepticism, in the simplest case, is to show that and why a certain kind of ampliative inference is rational—one which passes from finitely many observed examples of (and no observed counterexamples to) a natural pattern to the projection that the pattern extends indefinitely to unobserved cases. The problem is sometimes presented (though not actually by Hume³) as that of finding supplementary premises to render the type of inference in question *deductively* valid, and then to explain how such premises

3. Hume merely challenges his reader:

'But if you insist that the inference is made by a chain of reasoning, I desire you to produce that reasoning' (*An Enquiry Concerning Human Understanding*, Section 4, Part 2, 29).

might be justified. This, though, misrepresents matters. There is nothing general which we actually believe that will serve to transform an ampliative inductive inference into a deductively valid one. Consider—forgive the usual simple-minded schema—the inference from ‘All observed Fs are G’ to ‘All Fs are G’. The belief that this is—in the right kind of context, and subject to the appropriate controls—a reasonable inference is based on the thesis (what used to be called the Uniformity of Nature) that the world abounds in natural regularities. But that thesis does not provide a premise which, conjoined with the datum that all observed Fs are G, will entail the conclusion that all Fs are G. That argument is *still* deductively invalid. The role of the Uniformity Thesis is rather to provide an *informational setting* in which the observed pattern of co-occurrence between Fs and Gs defeasibly warrants generalisation. The contention of the inductive sceptic is then that there is an implicit circularity in our procedures. Without the collateral information of the Uniformity Thesis, no inductive inference, even the very simplest, is reasonable. But the only way in which the Uniformity Thesis might itself be justified is by inductive inference. Or so the sceptical thought runs.

It’s not often observed that this pattern of scepticism generalises—that essentially this form of argument may be put to the service of scepticism about each of, for example, the material world, other minds and the past. Let P be any proposition purporting to express a routine observation about my local perceptible environment—say: that I have two hands—and consider the following trinity:

- I My current experience is in all respects as if P
- II P
- III There is a material world

Here, the sceptical thought is, proposition I typifies the best possible evidence anyone can have for P—evidence, plausibly, such that if it and its ilk are not sufficient evidence for claims about the material world, then nothing is—and yet, as in the case of induction, movement from I to II is ampliative: the inference is a defeasible one. Moreover, so the sceptical thought contends, the evidential bearing of I on II is not unconditional: the warrant provided by I for II is, as I’ve expressed the matter elsewhere,

information-dependent.⁴ And paramount among the pieces of information that have to be in place in order for the move from I to II to be warranted is III: that there is a material world in the first place (whose characteristics, at least at the level of description typified by P, are representable, and normally successfully represented, in sense experience.) But the only foreseeable way of acquiring a warrant for III, so the argument goes, would be to infer it from a warranted belief of the kind typified by P. So, again, there is a vicious circle: it is if only I can get a warrant for a specific proposition about it that I can acquire a warranted belief that there is a material world, yet it is only if the latter is already warranted and part of my collateral information that I can draw on my experience to provide warrant for specific beliefs about it.

It's obvious enough that the same pattern of sceptical argument—I'll call it the I-II-III argument—can be enveloped around each of

I	(Where X is distinct from oneself) X's behaviour and physical condition are in all respects as if she was in mental state M	I	It seems to me that I remember it being the case that P yesterday
II	X in mental state M	II	It was the case that P yesterday
III	There are minds besides my own	III	The world did not come into being today replete with apparent traces of a more extended history

And, although the case does not schematise quite so succinctly, the paradox also afflicts so-called *abductive* inference, or inference to the best explanation (when realistically conceived). Roughly, inferring from a body of attested empirical generalisations to a theory which purportedly depicts the underlying causes of their holding is justified, it will be contended, only in the context of the collateral information that there is an appropriate underlying realm of causes in the first place—yet that is something which in turn could only be known by inference from prior knowledge of the truth of particular such theories.

4. Wright [2002].

To generalise. A version of this paradox will be available whenever we are persuadable (at least temporarily) that the ultimate justification for one kind of claim—a type-II proposition—rests upon ampliative inference from information of another sort—type-I propositions. In any such case the warrantability of the inference will arguably depend upon the presupposition that there is indeed a tract of reality suitable to confer truth on type-II propositions in the first place, a domain whose details are, in the best case, broadly reflected in type-I information. *A fortiori*, it will depend on the first component of that: that a domain of fact which type-II propositions are distinctively apt to describe so much as exists. That is the relevant type-III proposition—a proposition of sufficient generality to be entailed by any type-II proposition. The schematic form of the resulting sceptical argument is then given by these five claims:

- (i) Type-II propositions can only be justified on the evidence of (by ampliative inference from) type-I propositions.
- (ii) The evidence provided by type-I propositions for type-II propositions is information-dependent, requiring (among other things) collateral warrant for a type-III proposition.
- (iii) So: type-III propositions cannot be warranted by transmission of evidence provided by type-I propositions for type-II propositions across a type-II to type-III entailment—rather it's only if one already has warrant for the type-III proposition that any type-II propositions can be justified in the first place.
- (iv) Type-III propositions cannot be warranted any other way.

If all four propositions are accepted, then type-III propositions are cornerstones for type-II propositions (thesis ii) which cannot themselves be warranted (theses iii and iv). So

- (v) There is no warrant for any type-II proposition.

No doubt the justificational architecture postulated by the I-II-III argument is contestable in some of its local applications but there seems to be no hope whatever for the thought that it might successfully be contested *everywhere*. The relevant structure seems to be implicit in the very idea of *cognitive locality*.

Cognitive locality is the circumstance that only a proper subset of the kinds of states of affairs which we are capable of conceptualising is directly available, at any given stage in our lives, to our awareness. So knowledge of, or warranted opinion concerning the remainder must ultimately be based on defeasible inference from materials of which we are so aware. As we observed, type-III propositions are implicitly in play whenever our best justification for the truth of propositions of one kind—propositions of one distinctive type of subject matter—consists in the assembly of information about something else. That's the architecture which I-II-III scepticism attempts to impose, with varying degrees of plausibility, on the justification of propositions about the material world, about the past, about other minds and on inductive justification. And wherever such is indeed the justificational architecture, it will be plausible that a type-III proposition will form part of the informational setting presupposed in order for the relevant transitions to rank as warranted. Putting the matter in the most abstract form: suppose it granted that the best justification we can have for a certain kind of proposition—P-propositions—consists in information of another kind—Q-propositions—such that no finite (consistent) set of Q-propositions entails any P-proposition.⁵ The use of P-propositions in accordance with this conception will then carry a double commitment: a commitment to there being true P-propositions—and hence truth-conferring states of affairs for them—at all, and a commitment to a reliable connection between the obtaining of such truth-conferrers and the truth of finite batches of appropriate Q-propositions. That is the broad shape of the commitment which surfaces in the specific instances:

that there is a material world, broadly in keeping with the way in which sense experience represents it;

that other people have minds, whose states are broadly in keeping with the way they behave;

that the world has an ancient history, broadly in keeping with presently available traces and apparent memories;

that there are laws of nature, broadly manifest in finitely observable regularities,

5. This way of putting the point requires, if it is to be fully general, that infinite conjunctions of Q-propositions do not count as Q-propositions.

where each first conjunct presents a type-III proposition, while the second conjunct effects the connection necessary for the favoured kind of evidence to have the force which we customarily attach to it.

Philosophers may argue about, and be more or less generous concerning, what should be regarded as cognitively local. Descartes, at least for the purposes of the project of the *Meditations*, was relatively miserly, restricting the cognitively local to what was available to his reason and to certain forms of psychological self-knowledge. Twentieth century direct realism, by contrast, in the spirit that informs John McDowell's *Mind and World* and Hilary Putnam's Dewey Lectures,⁶ has been much more generous. But however generous one wants to be, a bound will surely have to be placed on cognitive locality at some point. Concerning what lies beyond it, our options will then be to regard it either as lying beyond our ken altogether, or as accessible to us only via the kind of inferential routine which the I-II-III argument purports to show is viciously circular.

Notice once again, however, that the thrust of this second genre of sceptical argument is that an evidential justification for the cornerstone—the type-III proposition—cannot be *acquired*: the claim is that in order to arrive at such a justification, one would have first to accomplish a process of justification (for a type-II proposition) which would presuppose it. So again there is a lacuna between the most that is strictly accomplished by the sceptical argument—that evidence for a cornerstone cannot be acquired by any foreseeable justificatory process—and the claim, that we have no warrant for it, which is actually what is needed to elicit the catastrophic conclusion (that there is no warrant for any belief of the type-II in question).

If I am right that the two distinguished—Cartesian and Humean—forms of sceptical argument between them capture, in essentials, all that we have to worry about, then their common lacuna suggests a common strategy of response—what I will call the *unified strategy*. Suppose there is a type of rational warrant which one does not have to *do any specific evidential work* to earn: better, a type of rational warrant whose possession does not require the existence of evidence—in the broadest sense,

6. Putnam [1994].

encompassing both *a priori* and empirical considerations—for the truth of the warranted proposition. Call it *entitlement*. If I am entitled to accept P, then my doing so is beyond rational reproach even though I can point to no cognitive accomplishment in my life, whether empirical or *a priori*, inferential or non-inferential, whose upshot could reasonably be contended to be that I had come to know that P, or had succeeded in getting evidence justifying P. The sceptical arguments purport to show that the rejection of Cartesian scenarios, and the acceptance of type-III propositions, are both beyond warrant by such investigative accomplishment. If they were nevertheless entitlements, warranted without evidence—whether by my own work, or that of experts in my community, or that of my precursors—no sceptical conclusions need automatically follow. I would be entitled to discount the idea that my experience might be no more than a sustained lucid dream, and entitled to accept that there is a material world just as we ordinarily suppose. The cornerstones could warrantably remain in place, even though it was conceded that our right to leave them there was unsupported by evidence for their truth. And, that being so, it wouldn't matter if, just as the sceptical arguments contend, they are indeed cornerstones whose removal would be catastrophic.

The suggestion merely that there are such things as entitlements in this general sense—much less that they extend far enough to service the serious anti-sceptical mission called for by the unified strategy—may seem like wishful thinking. Still, my purpose here is to take it seriously, though a discussion on the present scale is bound to leave many loose ends. In what follows I'll try merely to outline a *prima facie* case for a number of different possible species of entitlement and review some of the salient obstacles and further issues. The overall upshot will be, I believe, a prospect of at least some partial successes, and a clearer sense of what it might take to execute the unified strategy right across the board and of its foreseeable limitations.

II

Belief and Acceptance. An issue that needs to be considered immediately is what exactly entitlement would be a warrant to *do*. It doesn't just go without saying that it would be warrant to

believe the proposition in question—there are issues about how ‘belief’ should be understood. One reason why it is easy to overlook the lacuna in the sceptical arguments is because it can seem impossible to understand how it can be rational to believe a proposition for which one has *absolutely no evidence*, whether empirical or *a priori*. That a warrant to believe that *someone else* is not currently undergoing a sustained lucid dream would have to be evidence-based seems absolutely compelling; how can it make a difference if the subject involved is oneself? Likewise, that it takes evidence to provide warrant for believing a particular (type-II) proposition about the material world seems incontestable—how can it make a difference if one simply escalates the generality of what is believed (up to a type-III proposition)? The idea of a non-evidential warrant to *believe* a proposition can easily impress as a kind of conceptual solecism.

I do not myself know whether the notion of belief *is* actually so tightly evidentially controlled as to underwrite that impression. But at this stage of our discussion, I think the best tactic with the point is to grant it and see where that leads. Let’s accordingly concede that entitlement will be best conceived as something other than a kind of warrant to believe. What else could it be? What is required is that there is a mode, or modes, of *acceptance* of a proposition which can be rational but which are not tantamount to believing it in the conceded central sense of ‘belief’.

It is plausible that we do have a notion—in fact a variety of notions—of this kind. We register such modes of acceptance in our ordinary thought and talk when we speak of someone as, in a particular situation, *acting on the assumption* that P—as, for example, when one’s manner of driving a car may be structured by the assumption that every other motorist one comes across is a dangerous fool—or as *taking it for granted* that P, as when the protagonists in a court of law are required to take it for granted that the prisoner is innocent until proved guilty, no matter what they actually believe on the matter—or as *trusting implicitly* that P, as perhaps in the matter of the reliability of travel directions or the time of day from strangers. In the same ilk, van Fraassen⁷ famously proposed a distinction between acceptance of an

7. van Fraassen [1980].

empirical theory and believing it to be true, arguing that evidence of a theory's empirical adequacy justifies one in doing no more than the former. Here I do not think it does justice to his intent if we equate acceptance of a theory with the very belief that it's empirically adequate—accepting a theory is rather a further thing which believing it to be empirically adequate is supposed to justify one in doing. What it justifies one in doing is, roughly, behaving in all—or very many—respects as one would do if one believed the theory to be true.

As a first approximation, then, we may propose the notion of acceptance of a proposition as a more general attitude than belief, including belief as a sub-case, which comes apart from belief in cases where one is warranted in acting on the assumption that P or taking it for granted that P or trusting that P for reasons that do not bear on the likely truth of P. Of course one may—sometimes irrationally—also believe P in such cases, in the sense implicit in a conviction that one knows that P. Successful sceptical arguments may then embarrass such convictions. The aim of the unified strategy, however, will be to show that such scepticism may prove to carry no challenge, nevertheless, to the corresponding *acceptances* and that warrant to accept—rather than to believe—cornerstone propositions may be enough to block the sceptical paradoxes that attend arguments to the effect there is no such thing as getting evidence to believe them.

I'll have plenty more to say about the notion of acceptance in what follows, and we will eventually converge on one particular kind of acceptance as the most germane to the purposes of the unified strategy. But to conclude this section, let me quickly respond to a fairly immediate concern about the strategy.⁸ How exactly does it promise to shore up the possibility of justified belief in type-II propositions? We are proposing to concede, after all, that we may indeed have no (evidentially) *justified belief* in type-III propositions—that maybe we can point to no cognitive accomplishment of which the effect is a reason to take it that they are more likely to be true than not—but countering that we may nevertheless be rationally entitled to accept them. But if standard closure principles govern justified belief, then the counter comes too late to do any good. Standard closure principles will have it

8. Urged on me by Stephen Schiffer.

that justified belief in a type-III proposition will be a necessary condition for justified belief in anything one knows to entail it. To surrender the former will therefore be to surrender justified belief in type-II propositions more or less across the board. Maybe an entitlement to accept them nonetheless can be salvaged. But the idea was to use entitlement to save justification, not to replace it.

The observation is well made. It teaches that a proponent of the unified strategy must indeed impose some qualification on standardly accepted closure principles. In particular, it cannot be that evidentially justified belief is closed under (known/justifiably believed) entailment. That is not so remarkable a concession once one notices that evidential relations themselves are not so closed.⁹ But if we let ‘warrant’ disjunctively cover both evidential justification *and* entitlement, it can still be that warrant, inclusively so understood, obeys closure principles suitable to do justice to our strong intuitive conviction that ‘justification’—pre-theoretically understood—should do so.

The specific concern about closure should be distinguished from a more general concern it illustrates: that once we admit mere entitled acceptances into the role of cornerstones, we are bound to risk ‘leaching’, as it were—an upwards seepage of mere entitlement into areas of belief which we prize as genuinely knowledgeable or justified. I’ll come back to this more general concern towards the end.

III

Strategic Entitlement. One initially promising-looking direction is illustrated by Reichenbach’s famous work on the problem of induction.¹⁰ Imagine Crusoe starving hungry on his desert island and totally unsuccessful in his attempts to find any animal or marine food sources. There are, however, plenty of luridly coloured fruits, of various kinds, all strange to him and none, so far as he can see, being eaten by any of the small number of

9. Which I take to be the minimal lesson of the kind of purported counterexample to closure of knowledge, or justification, originally pressed by Dretske [1970]. Whether or not I am justified in believing that the celebrated stripy animals are zebras but not in believing that they are not mules cleverly disguised as zebras, I unquestionably have the evidence of my eyes for the first and not for the second.

10. See especially Reichenbach [1938], §38.

seabirds that occasionally visit the island (there seem to be no avian land species there). In these circumstances, Crusoe may quite understandably feel that he has absolutely no reason to believe that any of the fruits are safe for consumption, much less nutritious. Nevertheless it's clear, assuming an interest in survival, that he is warranted in eating the fruit. Eating the fruit is, in game theoretical parlance, a *dominant* strategy. If the fruit is edible, he survives by eating it and will not otherwise do so; if the fruit is inedible, eating it will do him no good and may do him some harm—but the worst harm that it may do will be no worse (anyway, let's suppose he so views matters) than the harm of starvation. In all relevant possible futures, the mooted course of action either works out better than all alternatives or no worse than any alternative.¹¹

The outlined reasoning justifies a course of *action* that would also be justified by (evidence for the) belief that the fruit is edible. As remarked, though, there is in the circumstances described no evidence for that belief. However in order for the Reichenbachian train of thought to serve the present purpose—that of assisting the unified strategy—something *attitudinal* has to be elicitable from it. In particular, we want to disclose reason to *accept* type-III propositions even if the possibility of obtaining evidence for them is allowed to be foreclosed by the sceptical argument. So

11. Here are three of Reichenbach's own examples:

A blind man who has lost his way in the mountains feels a trail with his stick. He does not know where the path will lead him, or whether it may take him so close to the edge of a precipice that he will be plunged into the abyss. Yet he follows the path, groping his way step by step; for if there is any possibility of getting out of the wilderness, it is by feeling his way along the path (Reichenbach [1949], p. 482).

The man who makes inductive inferences may be compared to a fisherman who casts a net into an unknown part of the ocean—he does not know whether he will catch fish, but he knows that if he wants to catch fish he has to cast his net. Every inductive prediction is like casting a net into the ocean of the happenings of nature; we don't know whether we shall have a good catch, but we try, at least, and try by the help of the best means available (Reichenbach [1968], pp. 245–6).

An example will show the logical structure of our reasoning. A man may be suffering from a grave disease; the physician tells us: 'I do not know whether an operation will save the man. But if there *is* any remedy, it is an operation.' In such a case, the operation would be justified. Of course, it would be better to know that the operation will save the man; but, if we do not know this, the knowledge formulated in the statement of the physician is a sufficient justification. If we cannot realise the sufficient conditions of success, we shall at least realise the necessary conditions of success. If we were able to show that the inductive inference is a necessary condition of success, it would be justified; such a proof would satisfy any demands which may be raised about the justification of induction (Reichenbach [1938] p. 349).

what, if any, attitudinal pay-off is there from the game-theoretic style of reasoning? Is there, on any natural understanding, a warrant provided for Crusoe's *accepting* that the fruit is edible, and should we think of him as implicitly doing so when he goes ahead and eats the fruit?

It's certainly open to us to fix such a use of the word, of course. Acceptance will stand in the appropriate generic relationship to belief just provided we regard an agent as accepting a proposition in all cases where she acts in a way which, given other aspects of her attitudinal set, would be rationally explained by her believing that proposition. So on this proposal, Crusoe, if persuaded to eat the fruits by the (cogent) reasoning outlined, thereby (warrantedly) accepts that the strange fruits are nutritious. And in general, the things one accepts will be the things one behaves—at least to a certain extent¹²—as if one believed. In cases when the explanation of that behaviour is strategic, as in Crusoe's situation, rather than attributable to an agent's actually having evidence for the belief in question, we may then speak of a *mere* acceptance.

That's a possible linguistic proposal. But the resulting use of 'acceptance' may seem forced and psychologically artificial. Consider this example.¹³ You've just passed through airport security when an insurance company representative approaches you saying that, as a promotion, his company is offering free travel and accident insurance to every hundredth passenger entering the departures lounge, the only cost being that you leave him your postal address for further promotions. Suppose you are indifferent to that cost or even mildly interested to learn what the company has to offer. It then seems manifestly rational to accept the free policy, by reasoning directly analogous to Crusoe's. If the plane doesn't crash, you'll suffer no harm; and if it does crash, it will have been in your interest—in the extended sense in which one has an interest in the welfare of one's heirs after one's death—to have had the policy. Accepting the offer is therefore a dominant play. But it seems very strange to say that you thereby also accept—even if you do not believe—that the plane is going to crash, or even that you are acting on that assumption. No interesting attitudinal state would seem to be entrained.

12. More needs to be said about to what extent—see below.

13. Due to Stephen Schiffer.

There are however two salient differences between the examples. First,¹⁴ whereas you doubtless expect to complete your journey safely, Crusoe has no particular reason to expect that the fruit is *inedible*. If he had, it might still be rational for him to behave in just the same way—eating the fruit might still represent his only chance, though now, subjectively speaking, a reduced one. But then, as in the airport insurance case, it would also seem intuitively wrong to speak of him as *accepting* that the fruit was edible. Second, the airport example does not actually involve what it was proposed that acceptance should minimally involve, viz. the agent's acting in a way which, given other aspects of his attitudinal set, would be rationally explained by his believing the proposition. On natural assumptions, what would be rationally explained by your believing the plane is going to crash would be, not your acceptance of an offer of free insurance, but your refusing to board the plane.

Two revisions to the proposal are thus invited: first, that acceptance that P should require *absence of disbelief* that P: agents can be properly said to accept a proposition only when it is rationally *available* to them; that is, is consistent with what they believe. And second, acceptance of a proposition should require that an agent really does act in *all* respects as if they believed the proposition in question, and not merely in some restricted salient set of respects.

In fact, the second revision entails the first. If I actually disbelieve P, then that fact is inevitably going to impact on the explanation of various things I do, or would be willing to do; so my behaviour will necessarily not be in *all* respects as if I believed P. But the second revision is also too strong as formulated. For if the notion of acceptance is to be well conceived, there had better—of course—be *some* operational differences between an agent who (rationally) merely accepts a given proposition and one who (perhaps irrationally) believes it. And there will. If Crusoe believes the fruit is edible, then his mood, for example, as he eats it is likely to be very different from how it will be if he merely accepts that it is edible in the strategic kind of sense proposed. The matter needs a much more nuanced discussion than I have space to attempt here but the relevant basic point is

14. Observed in discussion by Mauricio Suarez.

that the explanatory parallels between belief and mere acceptance will be restricted to their role as *reasons* for further thought and actions. Ordinary intentional psychological thought routinely involves a number of other—non-rationalising—kinds of explanation. There are many kinds of response—emotions, attitudes and actions—which it views as characteristic or expressive of a given psychological state without being *rationalised* by it. The constitutive requirement on an agent's acceptance that P should be that she (be disposed to) accept the consequences of P and to behave, in so far as she behaves rationally, just in ways that would be practical-syllogistically *rationalised* by her actual desires and other beliefs/acceptances if in addition she were also to believe that P. Note that the second revision, appropriately qualified in this direction, is still going to be strong enough to entail the first.

There is much more to say but let's take stock. There *is* an attitude which can naturally be associated with (some instances of) Reichenbachian reasoning: roughly, that of *committing oneself to act on a certain assumption*. That is certainly an attitude to the content of the assumption in question. Of course the scope of the commitment can be qualified—it may apply just to action for one specific goal in one specific context—and its duration can be relatively ephemeral. Crusoe's commitment will endure just so long as he needs nutrition and no other possible way of getting it but eating the fruit obtains. What the Reichenbachian thought provides us with is one relatively clear paradigm of how such a commitment can be rational for reasons which do not impinge on the likelihood of the truth of the assumption in question. But it will seem more natural to describe such a rational commitment as involving an attitude of *acceptance* to the extent that its rationality generalises across a variety of situations and contexts and, in the limit, across situations and contexts in general. Beliefs of course can change. But so long as I have it that someone believes something, then *ceteris paribus* that piece of information goes into the explanatory machinery to which I may appeal in rationalising his actions in *any* context. 'Acceptance' will be most naturally reserved for an attitude with a similarly wide explanatory potential.

The foregoing illustrates the general point that, so long as it is insisted that rational belief is *per se* belief supported by

evidence, the unified strategy must ultimately rest on a developed philosophical psychology of an attitudinal state, or states, of acceptance which are belief-like, and capable of underwriting belief, yet contrast with belief. Very roughly, if we think of 'belief', in its core uses, as denoting a normatively constrained and normatively constraining state—a state identified by its 'in-' and 'out-rules', as it were: something *essentially rationally controlled* by evidence and *essentially rationally committal* to thought and action—then the general idea I am canvassing is that it will be necessary, in trying to make something of the notion of rational entitlement, to think in terms of attitudinal states which share much of the second ingredient—the element and style of commitments involved—with belief, but *not the first*. Of course it will do no harm to call states of both kinds 'beliefs'. But then the sceptical point—that certain of our cornerstone 'beliefs' seem to be essentially uncontrolled by any proper accumulation of evidence—will no longer carry an automatic critical impact. The question will be whether these 'beliefs' are properly viewed as subject to such controls in the first place or whether they are not instead examples of a species of attitudinal acceptance whose rationality, when it is rational, may be grounded differently.

Here in any case, to round off this section, is a first proposal about entitlement:

An agent X is *contextually strategically entitled* to accept P just in case

- (i) X has no sufficient reason to believe that P is untrue; and
- (ii) in the particular context and for its characteristic purposes, it is a dominant strategy for X to act—as far as the achievement of those purposes is concerned—as if he had a justified belief that P.

So then

A thinker X is *absolutely strategically entitled* to accept P just in case

- (i) X has no sufficient reason to believe that P is untrue; and
- (ii) in all contexts, it is a dominant strategy for X to act exactly as if he had a justified belief that P.

In these terms, the Reichenbachian contention about induction may be viewed as being that an acceptance of the Uniformity Thesis is an absolute strategic entitlement, and that we are accordingly justified in basic inductive inference in contexts in general. By contrast, Crusoe is merely contextually entitled to accept that the island fruits are edible. And the airline insurance example involves no, even merely contextual, strategic entitlement.

IV

How Much Can Strategic Entitlement Do? How far might strategic entitlement, roughly so characterised, go towards executing the unified strategy?

Let's review the core thought as it concerns induction. We can represent it like this:

- (a) We need, if we are to lead even secure, let alone happy and valuable lives, to be able to form reliable (conditional) expectations about the future—about what will happen (if so-and-so is the case.)
- (b) (i) If the Uniformity Thesis holds, inductive methods will be the most effective way of arriving at the true generalisations and theories which will support such reliable expectations.
- (ii) If the Uniformity Thesis fails, no methods will do any better than induction as a means for arriving at reliable expectations.

Therefore

- (c) The use of inductive methods is a dominant strategy for arriving at reliable expectations.
- (d) We have no reason to believe that the Uniformity Thesis fails.

Therefore

- (e) We are absolutely strategically entitled to accept the Uniformity Thesis and, hence, to accept that the world is inductively amenable.

If this reasoning is accepted, it immediately provides a modest fire-wall around inductive scepticism. Sure, it doesn't give us the

right to say that we *know* that nature will be continually inductively amenable. Indeed it appears to provide no reason for the subjective confidence which we undoubtedly repose in inductive method—someone who grasps and acts on the strategic reasoning could quite consistently, it seems, be as pessimistic about induction as you like. (However, I will return to qualify this, and the qualification will be important.) Still, what primarily seems disconcerting about the sceptical argument is the apparent implication that there is no rational basis for preferring the methodology of empirical science to divination of entrails or the tarot pack. Reichenbach complains about Hume that

he is not alarmed by his discovery; he does not realise that, if there is no escape from the dilemma pointed out by him, science might as well not be continued—there is no use for a system of predictions if it is nothing but a ridiculous self-delusion . . . if there is no justification for the inductive inference, the working procedure of science sinks to the level of a game . . .¹⁵

If the reasoning to (e) above is effective, then—in perfect accord with the unified strategy—it pre-empts this depressing prospect. Maybe we do not know that Nature is Uniform and have no genuine evidence for the likelihood of its continuing inductive amenability. But if the argument succeeds, we are absolutely strategically entitled to accept that things will so continue. The methodology of empirical science will have a rational authority, at least insofar as it rests on simple inductive inference, which divination of entrails and readings of the tarot pack cannot match.¹⁶

I won't here consider further whether (any version of) the Reichenbachian argument should indeed be accepted.¹⁷ One obvious point of vulnerability is claim (b)(ii): worlds in which Uniformity fails, one might suppose, would come in all sorts of chaotic varieties—how can we be sure, *a priori*, that there are

15. Reichenbach [1938], p. 346.

16. This much is not yet a response to the 'leaching' problem highlighted at the conclusion of Section II above.

17. The argument has of course been roundly criticised. But perhaps the most major concern about it is its seeming inability to address Goodman's 'New Riddle' (Goodman [1955]). Even if inductive generalisation of sampled evidence is a dominant strategy, that gives one no guidance about the proper description (green emeralds or grue emeralds?) of the pattern displayed by the sample.

none in which some non-inductive method of belief formation might not be predictively more successful?¹⁸ Our interest now, however, is in the question, with what degree of success we might expect to be able to wield the emergent notion of strategic entitlement against I-II-III scepticism in general—granted, for the sake of argument, that it may prove to carry some clout against inductive scepticism in the fashion Reichenbach hoped.

Is there any possibility that we might make out an absolute strategic entitlement to accept the type-III propositions earlier reviewed? Consider the case of perception and the material world. *Prima facie* it's straightforward to generate an analogue of the reasoning from (a) to (e). Thus:

- (a)* It is of paramount importance to us to find our way around the world, make use of its resources, avoid danger, and so on. If we are to do these things, we need to be able to form reliable beliefs about the locations and dispositions of material objects.
- (b)* (i) If the world is generally open to our perceptual faculties, ordinary observation will be the most effective way of forming such beliefs.
- (ii) If the world is not generally open to our perceptual faculties, no other capacities that we possess will fare any better.

Therefore

- (c)* Reliance on ordinary observation is a dominant strategy for arriving at reliable beliefs about the location and dispositions of material objects.
- (d)* We have no reason to believe that the world is not generally open to our perceptual faculties.

Therefore

- (e)* We are absolutely strategically entitled to accept that the world is generally open to our perceptual faculties.

And that, if allowed, would certainly be a finding worth having. But it's clear, on reflection, that it comes short as a response to material world scepticism—in both Cartesian and I-II-III

18. But *can* Uniformity fail altogether? Some of the relevant mathematical issues here are usefully outlined in A. W. Sudbury [1972].

varieties—even in the restricted (non conviction-justifying) way in which (e) responds to inductive scepticism. It comes short as a response to Cartesian scepticism because an entitlement to accept that the world is generally open to our perceptual faculties *when they are engaged* seems to sit quite comfortably alongside the worry that, right now, those faculties are not engaged—that right now I am suffering a persistent lucid dream, or that I am (since yesterday, say) a brain-in-a-vat. And it comes short as a response to I-II-III scepticism about the material world because more is needed, according to that form of scepticism, to facilitate the inference from the relevant kind of type-I proposition to the relevant kind of type-II proposition than is provided by the entitlement which the reasoning actually promises to provide.

To see the last point, reflect that in order justifiably to move from a claim of type-I, concerning how things currently seem according to my experience, to a claim of type-II, concerning characteristics of local material objects, we require, according to the sceptical argument, the collateral information that:

There is a material world, broadly in keeping with the way in which sense experience represents it.

This embeds two components: the *ontology* of the material world, and the *methodology* of reliance on sense-perception as a source of belief about it. And the Reichenbachian routine, (a)*–(e)*, bears in effect only on the second. Once it is granted that there is a material world at all, we get a strategic entitlement—if the routine succeeds—to take it that it is, broadly, open to our perceptual capacities. But the reasoning simply helps itself to the ontological component—that there is an external material world at all—from the start; it is a presupposition of its premise, (a)*. No strategic entitlement issues to accept that there is a material world: only, if there is one, to accept that our sense experience yields broadly reliable representations of it. We may foresee a similar shortcoming in the attempt to address I-II-III scepticism concerning other minds and the past by versions of the same routine.

Why doesn't the Reichenbachian reasoning fall short in a similar way (again, I am not taking a stand on whether it falls short in other ways) as a response to inductive scepticism? Because inductive scepticism, though an instance of I-II-III scepticism, precisely differs from the other examples in targeting the *second*

component—the methodological component—in the collateral information which it claims is necessary if the relevant form of ampliative inference is to be justified. It is not in doubt—in the standard dialectic with the inductive sceptic—that there is indeed a (spatio-temporally) extended tract of reality going beyond hitherto observed regularities. The question is: what vindicates inductive method as a way of forming beliefs about that tract of reality? And the Reichenbachian answer, crudely, is that either it works (as well as or) better than anything else or nothing works at all. What the reasoning from (a)* to (e)* makes plausible is that a *methodological* scepticism about perception would be as tractable by such considerations as inductive scepticism is (if that is so tractable at all). Likewise for a methodological scepticism about reliance on others' manifest behaviour and physical condition as a guide to their mental states, or reliance on memory as a guide to the past. But scepticism about the material world, other minds and the past is classically not methodological but ontological. And for this ontological scepticism, it appears, entitlement of strategy promises no cure. If the unified strategy can offer a cure, it will be by means of a different medicine.

V

Entitlement of Cognitive Project. A second, rather different species of entitlement is suggested by one tendency in Wittgenstein's remarks *On Certainty*. Here are two illustrative passages:

163. ... We check the story of Napoleon, but not whether all the reports about him are based on sense-deception, forgery and the like. For whenever we test anything, we are already presupposing something that is not tested ...

Compare:

337. One cannot make experiments if there are not some things that one does not doubt. But that does not mean that one takes certain presuppositions on trust. When I write a letter and post it, I take it for granted that it will arrive—I expect this.

If I make an experiment I do not doubt the existence of the apparatus before my eyes. I have plenty of doubts, but not *that*. If I do a calculation I believe, without any doubts, that the figures on the paper aren't switching of their own accord, and I also trust my memory the whole time, and trust it without reservation.

To take it that one has acquired a justification for a particular proposition by the appropriate exercise of certain appropriate cognitive capacities—perception, introspection, memory, or intellection, for instance—always involves various kinds of presupposition. These presuppositions will include the proper functioning of the relevant cognitive capacities, the suitability of the occasion and circumstances for their effective function, and indeed the integrity of the very concepts involved in the formulation of the issue in question. I take Wittgenstein's point in these admittedly not unequivocal passages to be that this is essential: one *cannot but* take certain such things for granted (though I am not sure how we should interpret his implied contrast between taking for granted and 'taking on trust'. More about trust shortly.)

That is not to deny that, if one chose, one could investigate (at least some of) the presuppositions involved in a particular case. I might go and have my eyesight checked, for example. But the point is that in proceeding to such an investigation, one would then be forced to make further presuppositions of the same general kinds (for instance, that my eyes are functioning properly now, when I read the oculist's report, perhaps with my new glasses on.) Wherever I get in position to claim justification for a proposition, I do so courtesy of specific presuppositions—about my own powers, and the prevailing circumstances, and my understanding of the issues involved—for which I will have no specific, earned evidence. This is a necessary truth. I may, in any particular case, set about gathering such evidence in turn—and *that* investigation may go badly, defeating the presuppositions that I originally made. But whether it does or doesn't go badly, it will have its own so far unfounded¹⁹ presuppositions. Again: whenever claimable cognitive achievement takes place, it does so in a context of *specific* presuppositions which are not themselves an expression of any cognitive achievement to date.²⁰

19. —*unbegründet* (*On Certainty* §253).

20. It's natural to rejoin that one may have inductive grounds for confidence in the present sound functioning of one's perceptual faculties. But in that case one relies on the evidence for the induction—on the *previous* sound functioning of one's perceptual faculties. So was that independently checked in a large number of cases? And even if so, is not one in any case now relying, without specific evidence, on one's memory of the outcome of the checks?

These presuppositions are not just one more kind of Wittgensteinian ‘hinge’ proposition as that term has come generally to be understood. Hinges, broadly speaking, are *standing certainties*, exportable from context to context. Whereas the present range of cases are particular to the investigative occasion: they are propositions like that my eyes are functioning properly *now*, that the things that I am *currently* perceiving have not been extensively disguised so as to conceal their true nature, etc.

A natural first reaction is that if this is right—if all claimable cognitive achievement rests on specific, ungrounded presuppositions—then we just have the materials for a new—third form of—sceptical paradox. The key thought in the new paradox would be a generalisation of part of the Cartesian sceptical routine about dreaming. Plausibly, our confidence in the things which we take ourselves to have verified in a particular context can rationally be no stronger than our confidence in the kind of context-specific presuppositions just remarked. Suppose I set myself to count the books on one of the shelves in my office and arrive at the answer, 26. Then the warrant thereby acquired for that answer can rationally be regarded as no stronger than the grounds I have for confidence that I counted correctly, that my senses and memory were accordingly functioning properly, that the books themselves were stable during the count and were not spontaneously popping into and out of existence unnoticed by me, etc. Yet I will have done nothing—we may suppose—to justify my confidence in all these specific presuppositions. So how have I achieved any genuine warrant at all?

Here is a possible line of reply. If there is *no such thing* as a process of warrant acquisition for each of whose specific presuppositions warrant has already been earned, it should not be reckoned to be part of the proper concept of an acquired warrant that it somehow aspire to this—incoherent—ideal. Rather, we should view each and every cognitive project as irreducibly involving elements of adventure—I have, as it were, to *take a risk* on the reliability of my senses, the conduciveness of the circumstances, etc., much as I take a risk on the continuing reliability of the steering, and the stability of the road surface every time I ride my bicycle. For as soon as I grant that I ought—ideally—to check the presuppositions of a project, even in a context in which there is no particular reason

for concern about them, then I should agree *pari passu* that I ought in turn to check the presuppositions of the check—which is one more project after all—and so on indefinitely, unless at some point I can foresee arriving at presuppositions all of which are somehow safer than those of the initial project. If not, then there will be no principled stopping point to the process of checking: the quest for security will be endless, and therefore useless. And if that is the situation, then the right response—the reply will continue—is not to conclude that the acquisition of genuine warrant is impossible, but rather to insist that it does not require this elusive kind of security. Rather, warrant is acquired whenever investigation is undertaken in a *fully responsible* manner, and what the paradox shows is that full epistemic responsibility cannot, *per impossibile*, involve an investigation of every presupposition whose falsity would defeat the claim to have acquired a warrant. (Suggestion: the correct principle is not that any acquired warrant is no stronger than the weakest of one's independently acquired reasons to accept each of its presuppositions. It is, rather, that it is no stronger than the warrant for any of the presuppositions about which there is some *specific antecedent reason* to entertain a misgiving.)

This line of reply concedes that the best sceptical arguments have something to teach us—that the limits of justification they bring out are genuine and essential—but then replies that, just for that reason, cognitive achievement must be reckoned to take place *within such limits*. The attempt to surpass them would result not in an increase in rigour or solidity but merely in cognitive paralysis.

Let me try to harness these ideas to a definite proposal about entitlement. First (to tidy up a bit) a definition: let us say that

P is a *presupposition* of a particular cognitive project if to doubt P (in advance) would rationally commit one to doubting the significance or competence of the project.

Then the relevant kind of entitlement—an entitlement of cognitive project—may be proposed to be any presupposition of a cognitive project meeting the following additional two conditions:

- (i) We have no sufficient reason to believe that P is untrue

and

- (ii) The attempt to justify P would involve further presuppositions in turn of no more secure a prior standing . . . and so on without limit; so that someone pursuing the relevant enquiry who accepted that there is nevertheless an onus to justify P would implicitly undertake a commitment to an infinite regress of justificatory projects, each concerned to vindicate the presuppositions of its predecessor.

No doubt that will stand refinement, but the general *motif* is clear enough. If a cognitive project is indispensable, or anyway sufficiently valuable to us—in particular, if its failure would at least be no worse than the costs of not executing it, and its success would be better—and if the attempt to vindicate (some of) its presuppositions would raise presuppositions of its own of no more secure an antecedent status, and so on *ad infinitum*, then we are entitled to—may help ourselves to, take for granted—the original presuppositions without specific evidence in their favour. More generally, wherever we need to carry through a type of project, or anyway cannot lose and may gain by doing so, and where we cannot satisfy ourselves that the presuppositions of a successful execution are met except at the cost of making further presuppositions whose status is no more secure, we should—are rationally entitled to—just go ahead and trust that the former are met.

I said ‘trust that’ and not merely ‘act on the assumption that’. Here is the place to register a very important gloss on the understanding of ‘acceptance’ needed by the unified strategy. Earlier, in discussion of the Reichenbachian approach, it was suggested that the kind of acceptance which is motivated by an (absolute) strategic entitlement would be consistent with agnosticism, even pessimism about the truth of the supposition in question. That, of course, if correct, limits the power of strategic entitlement as a response to inductive scepticism, since there is no question but that ordinary inductive thought involves, more than a strategic acceptance of the inductive amenability of the world, an implicit trust in it. Do we do better in this respect with entitlement of cognitive project? Is it an entitlement to *trust*?

Suppose it is less. In that case, to appreciate my entitlement to accept that my sensory apparatus, for example, is right now

generally sound will be fully consistent with my taking an agnostic or sceptical view about the matter. But seems impossible to square agnosticism, say, about that with a conviction of the truth of the ordinary day-to-day things I routinely take myself to verify by perceptual means. I cannot *rationally* form the belief that it is currently blowing a gale and snowing outside on the basis of my present visual and auditory experience while being simultaneously agnostic, let alone sceptical, about the credentials of that experience. Sure, I can decide what beliefs it would be *appropriate* to form on the assumption that my sensory apparatus is currently sound, but I will not, if rational, be able to form those beliefs while I am open-minded—so unpersuaded—whether it *is* sound. To choose to act on an assumption is—extensionally—to choose to act in ways that would be rationalised by believing it. But this chosen range of action cannot, for a rational subject, extend to the formation of the beliefs that would be appropriate if, more, one trusted that the assumption was true. Since believing in general is not purely voluntary but is controlled by reasoning and evidence, it is not a rational option for someone who is sceptical or agnostic about the pedigree of the relevant evidence, or the character of the reasoning involved.

So much is indeed implicit in the very characterisation I gave of a presupposition of a cognitive project: something doubt about which would rationally commit one to doubting the significance or competence of the project. Since one will not (rationally) believe anything on the evidence afforded by carrying through a project of whose significance or competence one is unpersuaded, it follows immediately that if acceptance of such a presupposition is to be capable of underwriting rational belief in the things to which execution of the project leads, it has to be an attitude which *excludes doubt*. If there is entitlement of cognitive project, it has to be an entitlement not merely to act on the assumption that suitable presuppositions hold good, but to place trust in their doing so.

The same, indeed, must hold for absolute strategic entitlement too. More carefully, it must hold for absolute strategic entitlement to any supposition which is to underwrite a policy of *belief formation* (rather than merely non-doxastic forms of action). In particular, if a strategic acceptance of the Uniformity of Nature is to ground specific inductively formed *expectations*—contrast:

working hypotheses—then that acceptance cannot be exhausted by the decision merely to act on the assumption of Uniformity in all contexts. For again, free action on an assumption will—in a rational subject who is uncommitted to its truth—inevitably stop short of the formation of the specific beliefs which holding it to be true would mandate. So our discussion of Reichenbach and strategic entitlement needs a crucial amendment. If there is a strategic entitlement to a policy of forming beliefs inductively, it must be an entitlement to *trust* that the world is so constituted that such a policy will, by and large, often enough, be successful. A strategic entitlement to accept the Uniformity of Nature and Crusoe's strategic entitlement to accept that the island fruits are edible differ in just that respect.

This is the point of convergence I promised in Section II. Acceptance, for the purposes of the unified strategy, is—or has to involve—*trust*. 'Warrant for nothing' is entitlement to trust. It is in the nature of trust that it gets by with little or no evidence. That is exactly how it contrasts with belief proper, and it is not *per se* irrational on account of the contrast. Entitlement is rational trust.

VI

How Much Can Entitlement of Cognitive Project Do? We already touched on one striking prospective capture when we noted that the presuppositions of a given cognitive project will characteristically include the proper functioning of the cognitive capacities which need to be engaged in pursuing it, the suitability of the attendant circumstances for their effective function, and indeed the integrity of the very concepts involved in the formulation of the project to hand. Since this goes for any cognitive project, there are bound to be presuppositions falling within these same three broad categories which occur at the next level up—if one were to set out to confirm the presuppositions of the original project—and so on indefinitely. So, while the details need thinking through, there seems every prospect that some presuppositions of at least these three kinds will meet the defining conditions on entitlement of cognitive project. It would follow in particular—provided the very idea of entitlement of project is in good standing—that in all circumstance where there is no specific reason to think otherwise,

we are each of us entitled to take it, without special investigative work, that our basic cognitive faculties are functioning properly in circumstances broadly conducive to their successful operation. If so, that immediately empowers us to dismiss the various scenarios of cognitive dislocation and disablement—dreams, sustained hallucination, envatment and so on—which are the stock-in-trade of Cartesian scepticism. That, for instance, I am not right now dreaming is a presupposition in the sense defined of any cognitive project involving perceptual interaction with the world, and a presupposition, moreover, which I have, right now, no reason to suppose unsatisfied and of which any effective investigation by me would involve the same presupposition over again. That indeed was the triumphant thrust of the sceptical routine we reviewed at the start: that there is no evidentially justifying the claim that I am not right now dreaming. But under the aegis of entitlement of cognitive project, that routine is tamed to issue in the benign conclusion that I am rationally entitled to take the falsity of the dreaming hypothesis on trust in any broadly empirical cognitive project; so the Cartesian sceptical argument, which depends on my having no good reason to discount it, is nipped in the bud.

This is a good result, it goes without saying, only if it is selective—only if the entitlements generated turn out to be cornerstones of our actual ways of thinking about and investigating the world and do not extend to all manner of (what we would regard as) bizarre and irrational prejudices. As a test case, suppose I undertake a project is to predict the winners in tomorrow's card at Newmarket by rolling a pair of dice for each runner in the afternoon's races and seeing which get the highest scores. Clearly it is a presupposition of this project that the method in question has some effectiveness. What prevents that presupposition becoming an entitlement?

The obvious answer is that clause (i) is unsatisfied—there is every reason to doubt that the method in question is effective. But that is not the fundamental point. The fundamental point is that (as we know) it would be straightforward to gather no end of empirical evidence to discredit the dice-rolling method. And this would not be possible if the various presuppositions of such evidence-gathering in turn were of 'no more secure a prior standing' than the dice-rolling method. If they were of no more

secure a prior standing, we'd have to admit to a stand-off and suspend judgement. So the very discreditability of the method entails that clause (ii) is unsatisfied.

A doubt now comes into focus, however. There is no entitlement to trust in the dice-rolling method because it is a method for assessing statements which allow of independent assessment by more basic means, whose reliability is of more secure prior standing. What of a case where that feature is missing? Suppose I postulate a tract of reality—it might be the realm of non-actual possible worlds as conceived by Lewis—which is spatio-temporally insulated from the domain of our usual empirical knowledge, and a special faculty—as it may be, our non-inferential 'modal intuition'—whose operation is supposed to allow us to gather knowledge about it. Do I have an entitlement of cognitive project to trust the (alleged) faculty on any particular occasion? If not, why not? After all, I have—in the nature of the case, since I cannot compare its deliverances with the facts, independently ascertained—no reason to believe that it is unreliable (so long as its prompting are consistent); and any attempt to check on its functioning will presumably perforce involve further modal intuition, 'of no more secure a prior standing'. But do we want Lewis's views about the nature of modality—making no judgement about their independent merit—to turn out to be a matter of rational entitlement in any case? If not, what blocks them doing so?

The example highlights something vital about the limitations of this genre of entitlement. It may very well prove to be the case that a trust in the reliability of basic modal intuition—in our primitive, non-inferential impressions of modal validity and invalidity—turns out to be a matter of entitlement of cognitive project. What is not an entitlement—or not this kind of entitlement anyway—is the specific Lewisian metaphysics, or *any* specific metaphysics, of the nature of modal reality, any specific conception of the kinds of states of affairs which make modal claims true or false. We may prove to be entitled to trust, in any particular cognitive project involving modal judgement, that those of our faculties which are essentially involved in such judgement are functioning properly in circumstances broadly conducive to their effective function. But we are not thereby entitled to any particular conception of the nature of modal facts.

The point, generalised, is that entitlement of cognitive project fares no better than strategic entitlement as a response to I-II-III scepticism in general, and falls short in a similar way. Type-III propositions—that there is a material world, that there are other minds, that the world has an extended history—are indeed presuppositions of our enquiries in the sense defined. But they are not entitlements of cognitive project as characterised, since they fail to meet condition (ii). The problem with type-III propositions is not that—like ‘my visual system is functioning properly on this occasion’—to accept that there is an onus to justify them in any particular context in which they are presuppositional would—plausibly—be to accept an infinite regress of similar justificatory obligations. Rather, it is that, failing some independent response to the sceptical argument, one has no idea how to justify them at all. Entitlement of cognitive project does not, any more than strategic entitlement, extend to matters of ontology. Once granted a certain conception of certain of our cognitive powers and the nature of their sphere of operation, we may be able to appeal to this kind of entitlement to make a case that, for the purposes of any particular enquiry of the relevant kind, we are entitled to take it on trust that those powers are functioning effectively in conducive circumstances. But if so, this congenial finding comes within a context in which the broad nature of the powers in question and the character of their subject matter to which they are sensitive is not in question. Once those matters do come into question, it is hard to see that anything so far said promises much in the way of answers.

VII

Entitlement of Rational Deliberation. Strategic entitlement and entitlement of cognitive project both allow that it is rational to place trust, without evidence, in two kinds of presupposition of pure *enquiry*—those whose acceptance generates a dominant policy in relation to the goals of (a particular kind of) enquiry, and those which believing the results of an enquiry rationally requires us not to doubt, yet which are beyond vindication by evidence except at the cost of further presuppositions of the same kind (or more generally, further presuppositions which are no more

secure). In contrast, the third genre of entitlement to be canvassed here is anchored in the constitutive requirements of rational *action*.

The generic thought is that since rational agency is nothing we can opt out of, we are entitled to place trust in whatever (we have no evidence against and which) needs to be true if rational decision-making is to be feasible and effective. More carefully, say that P is a *general presupposition* of rational deliberation just in case it may be recognised *a priori* that a soundly based—justified and correct—decision on the respective merits of alternative courses of action open in a particular context is possible only if P holds good in that context—so that an agent who found herself in possession of reason to regard P as failing in a particular deliberative context would be bound to regard herself as—if only temporarily—incapacitated from rational decision-making. The proposal is, then, that an agent has an entitlement to place trust in any of the general presuppositions of rational deliberation which she has no reason to regard as failing in her particular deliberative context.

She is so entitled because the need to take decisions will, time and again, trump whatever may be the limited possibilities—especially in the light of sceptical argument—for gathering positive evidence that the general presuppositions hold good in the particular context, and because—as a rational agent—her decisions have to be informed by *reasoned beliefs* about what is for the best. Since such beliefs will be possible for her only in a context in which she has trust that what she knows to be necessary conditions for their being soundly arrived at are met, only a thinker who has such trust can be a rational agent.

Two such species of necessary conditions, hence entitlements, are worth remarking. Deliberating what to do involves consideration of alternatives, in the light of one's wishes and aims. So one general presupposition of rational deliberation is that one has sufficient self-knowledge to identify those of one's wishes and aims which are relevant to the decision at hand. Psychological self-knowledge, to that limited extent, is an entitlement of rational deliberation.

A second example emerges from the observation that sound assessment of alternative courses of action requires, *inter alia*, knowing or justifiably believing a range of (open subjunctive) conditionals which variously define what may be expected to occur

if such-and-such a direction is pursued. And such knowledge inevitably involves reviewing what support is provided for them by various relevant kinds of known generalisation: it is reasonable to believe that if I were to perform an action of such-and-such a kind, such-and-such a situation would (probably) result because it always (or usually) does, or because, although there is no established pattern, there is theoretical reason, backed by other generalisations, for that expectation. In general, ordinary rational deliberation is possible only for someone who rationally believes in certain relevant empirical generalisations. And it is well grounded only when suitable such generalisations are true.

Consider then the proposition that nature displays sufficiently many inductively and abductively ascertainable regularities to make the prosecution of those methods worthwhile. We have no reason to disbelieve this. And if it is not true, then we lose the means necessary to select the subjunctive conditionals that are needed in practical deliberation, and practical deliberation itself becomes paralysed. By the proposed notion of entitlement, we are therefore in position rationally to accept that nature displays sufficiently many inductively/abductively ascertainable regularities to make the prosecution of those methods worthwhile. But that is enough to ensure the rationality of employing those methods.

In more detail:

We cannot function effectively as deliberative agents unless we presuppose that there is a wealth of correct subjunctive conditionals.

Subjunctive conditionals are correct in virtue of nomic regularities.

So we cannot function effectively as deliberative agents unless we presuppose that there are sufficient nomic regularities to sustain a wealth of correct subjunctive conditionals.

Since there is no reason to doubt that there is such a sufficiency of nomic regularities, the supposition that there are is accordingly an entitlement by the proposal. But now suppose it may also be shown that

If a nomic regularity obtains and there are accessible grounds for believing in it at all, then here are broadly inductive, or abductive grounds for doing so—grounds that belong with the methodology of the developed empirical sciences.

Nomic regularities which make an observable difference, in other words, are symptomatised by the availability of scientific evidence that they obtain. It would follow that, in the context of our need to select correct subjunctive conditionals, and hence to find true regularities to sustain them, there is a better type of ground to rely on than the broadly scientific—say an *M-ground*—only if two conditions are met: first, that it likewise follows from the obtaining of a nomic regularity that, if warrantably believable at all, there will be M-grounds for believing it; and second, that there is less chance of *rogue* M-grounds—M-grounds which indicate a nomic regularity where there is none—than of rogue inductive or abductive grounds. But we know of no such type of ground.

And that, it would seem, is enough. Under the rubric proposed, we have an entitlement of rational deliberation to trust that there are many nomic regularities. Any nomic regularity that has an effect on the observable is symptomatised by the availability of inductive or less direct empirical evidence. We know of no more reliable symptom. So reliance on inductive and abductive methods is the best we can knowledgeably do in pursuit of purposes which are essential to rational agency itself, and thus unavoidable.

This attempt to capture induction as an entitlement is somewhat distant from the strategic route. In particular, there is no need for the lemma that inductive inference is a dominant strategy. The essential thought is merely, that the truth of the Uniformity Thesis is a general presupposition of rational deliberation and that, absent evidence to doubt it and knowing of no superior way in general to gain access to the regularities in whose existence we are thereby entitled to trust, it is rational to rely on inductive and abductive methods in doing so.

VIII

Entitlements of Substance? If, as just argued, entitlement of rational deliberation can be made to cover inductive and abductive inference, then it ought to allow extension to an acceptance of (records concerning) a substantial past. For it is the past that offers the evidence which those methods require. So at least in context of rational deliberation, material is promised to

address I-II-III scepticism in two of the originally troublesome areas—empirical generalisations and statements concerning the past. But we still have the problem that has beset us throughout: the prospects of making out that our acceptance of an external material world and the existence of other minds are entitlements seem no better in this context than they transpired to be under the aegis of entitlement of strategy and of cognitive project. Maybe some subtle philosophy can disclose otherwise but it is no obvious presupposition of rational deliberation to conceive of the stuff of the world as matter or to acknowledge the existence of minds besides one's own. Can these—our most fundamental conceptions of the substance of the world—be made out to be matters of entitlement? I have no definite argument to offer for either answer, but must be content merely to indicate a direction by which a (partial) answer might be found.

As I stressed earlier, we may avoid particular versions of the I-II-III argument by arguing for a rejection of the justificational architecture which it presupposes—with perceptual claims, perhaps, in pole position for the attempt. But if this is to be a *globally* successful tactic, then we will have to do nothing less than so fashion our thinking that it *nowhere* traffics in propositions related as type-I propositions and type-II propositions. *None* of the thoughts we think must be such that their truth-makers are beyond our direct cognition, so that we are forced to rely on finite and accessible putative *symptoms* of their obtaining.

Could there be such a way of thinking? Earlier, it was suggested that there could not; that an unavoidable and unacceptable casualty of any such scheme of thought would be the thinker's conception of her own *cognitive locality*—the idea of a range of states of affairs and events existing beyond the bounds of her own direct awareness. Globally to avoid the justificational architecture presupposed by I-II-III scepticism would be to forgo all conception of oneself as having position in a world extending, perhaps infinitely, beyond one's cognitive horizon. In particular, it would be to surrender all conception of our own specific situation within a broader objective world extending *spatially and temporally* beyond us. However it is, of course, a real and crucial question—not to be addressed dogmatically—whether there could be any coherent such system of thought.

I shall not here try to develop an argument that there could not. Certainly, all *our* actual thought and activity is organised under the aegis of a distinction between states of affairs accessible to us at our own cognitive station and others that lie beyond, and it is difficult to form any clear concept of how things might be otherwise. There is however a well-known train of thought, popularly understood as Kantian and given body by Strawson's classic discussion in Chapter 2 of *Individuals* and the memorable critique of Strawson by Gareth Evans,²¹ which argues, in effect, that cognitive locality goes with the very idea of our experience as being of an *objective* world, of a reality that stands independent of it. More specifically, it is only via a conception of the possibility of states of affairs and processes occurring *unperceived* that sense can be given to the idea that experience informs us of a reality not of its making. But that conception calls in turn for a conception of a way, or ways, in which states of affairs and processes can elude the awareness of a thinker, which—according to the Kantian train of thought—in turn necessarily involves some *dimension* of variation of locality—the idea of a situation obtaining, in the most abstract sense, *elsewhere*—and hence a conception of that dimension of variation. And now *that* conception in turn arguably demands some notion of the make-up—substance—of a state of affairs suitable to allow it to be situated 'elsewhere'.

Of course, that is all—to put the matter kindly—somewhat promissory. And it impresses as ambitious to hope that our specific conceptions of space (and time) and matter might somehow precipitate themselves out of this direction of enquiry as transcendently imposed by the very idea of objective experience. But something less specific might: it does not seem altogether fanciful that a developed (Kantian or Strawsonian) metaphysics might teach us that to operate any scheme of thought rich enough to recognise objective experience—rich enough to allow for experience of states of affairs whose existence is constitutively independent of experience—must involve a grasp of the idea of particular states, events and processes existing outside the thinker's cognitive locality, and hence some conception of dimension(s) of locality and an appropriately co-

21. P. F. Strawson [1959] and Gareth Evans [1980].

ordinated conception of substance. If so, then the mere conception of ourselves as capable of experience *of* a world cannot escape *some* conception of substance: of the nature of what fundamentally constitutes the kinds of states of affairs that can be situated ‘elsewhere’. Assuming that conceiving of experience as objective is somehow independently mandated—of course, that raises further major issues—a somewhat minimal notion of entitlement of substance might then emerge: since some conception of one’s cognitive locality and of the substance of states of affairs that are elsewhere is essential to any objective conception of experience—and since (suppose) so conceiving of experience is independently warranted or unavoidable—a thinker is entitled to the basic ontology involved in an otherwise coherent conception of what kind of thing might obtain at other localities.

Notice that this still falls short of requiring a *specific* ontology. Unless there are independent objections to any other ways of realising the general shape—or unless there are no such other ways—it merely gives us a *permission* for matter and space. And it says, so far, nothing about mind.

IX

Concluding Reflections. Let me review the main lines of the prospectus I’ve tried to motivate and highlight some of the matters that remain for further work.

‘Warrant for nothing’ is a nice phrase, but is entitlement—at least in the guise that has emerged—happily captioned by it? After all, there has been no suggestion that one is justified, by default and without evidence for their truth, in holding to certain *beliefs*. Rather, the proposal is that the idea of entitlement is best approached in terms of a more generic kind of propositional attitude—the kind of attitude signalled, at least in certain kinds of use, by the phrases ‘acting on the assumption that’, ‘taking it for granted that’ and ‘trusting that’. We can think of both belief and other attitudes of this kind as sub-species of a more general notion of *acceptance*, unified by analogies and overlap in what they respectively require of a rational agent who, in one or another way, accepts a given proposition. I have not here tried to fill out the detail of the analogies and overlap.

I have suggested that the relevant mode of acceptance for the purposes of the unified strategy is trust. It is in the nature of trust that it may be placed, without stigma, in things for which one has no evidence. But it is not *per se* ‘for nothing’ if that is understood as ‘normatively unconstrained’. Trusting without evidence can still be rational or not. Entitlements are *warrants to trust*, supported in the kinds of ways we have been reviewing. The basic respect in which ‘warrant for nothing’ is apt as a caption for entitlement is that recognising the rationality of trusting that P need involve none of the work—empirical or *a priori*—that would have to go into the accumulation, perhaps *per impossibile*, of evidence for believing P. Counting both entitlements and evidential justifications as types of warrant, then, entitlement is not, maybe, a warrant costing nothing at all but it is at least a warrant costing nothing of the kind that would be involved in getting evidence for the truth of the proposition in question.

That said, though, there is a further element in the way the caption is naturally understood. The question arises whether, in order to enjoy an entitlement to a particular proposition, one has *oneself* to accomplish the demonstration that there is such warrant—to recognise oneself that the case is one where trust is rational. In normal circumstances, and putting to one side special issues concerning testimony, a thinker’s knowing or justifiably believing something requires that she herself have evidence for it sufficient to constitute knowledge or justified belief. By contrast we do not, at least in a wide class of cases, demand that before a thinker can justifiably infer in accordance with a principle of inference, she must herself have accomplished a justification for the use of that rule. At least in cases where a valid pattern of inference demands no special training but is followed by the ‘light of natural reason’, we will naturally credit a thinker with warrant to proceed as she does, even if she has given no explicit thought to that way of proceeding and would not have the slightest idea how to answer if a request for justification was made. If entitlement stands comparison in this matter with justified belief rather than warranted inference, then—in the present state of our understanding of these issues—no-one yet has ever had much in the way of entitlements. Roll on the day when we get these things straightened out, and can at last get some entitlement to our cornerstones, fend off scepticism and start accumulating some

knowledge! Clearly this is yet another issue for further attention, but entitlement had better prove to be 'for nothing' in this additional sense too—had better be comparable to rights of basic inference, as it were—if any but a few philosophers are to benefit from a vindication of the notion. The matter is deep and the comparison with basic logic suggestive.

I have outlined a case for three kinds of entitlement—strategic entitlements, entitlements of cognitive project, and entitlements of rational deliberation—and have gestured, in the most promissory and indefinite way, at the possibility of—and need for—a fourth, entitlement of substance. Entitlement of cognitive project seems to promise well in addressing the challenge of Cartesian scepticism, or any variety of scepticism that works by trying to dislodge a cornerstone of our intellectual or cognitive competence. Entitlements of strategy, and of rational deliberation, promise to be of use in addressing the challenge of inductive—and, more generally, methodological—scepticism, and also scepticism concerning the reliability of the various kinds of cognitive faculty that enter into the ways we form beliefs about subject matters of all kinds—*provided* we are granted an unchallenged conception of such subject matters. But as far as I have been able to see, these three kinds of entitlement fall short of the materials needed for a complete execution of the unified strategy. Humean—I-II-III—scepticism about the material world, and about other minds, *does* challenge our conception of the kinds of subject matter which, at the most general categorical level, the world puts up for our consideration. An entitlement of substance—an entitlement to have a view about the most basic categories of stuff and thing the world contains—is what it would take to close this gap. But we have glimpsed the merest outline of a recipe for the beginnings of a case there are any such entitlements.

The situation has to raise a concern about the ability of the unified strategy to fend off an unwelcome pluralism. It is, to be sure, an essential feature of the notion of entitlement that it is a matter of *rational* trust—and that's sufficient safeguard, for all but the most thoroughgoing sceptic about rationality, that not just any old trustings will do. But the point remains: if the most favourable light that can be cast on our acceptance of a material world, or other minds, consists in argument that our very

rationality means we have to have *some* such commitments—for instance, that any system of rational objective thought has to incorporate *some* conception of the kind of stuff that inhabits other cognitive localities—then we seem to have no claim to the objective correctness of the most fundamental categories of substance that we actually employ. More, there will be no obstacle in principle to the idea of alternative, equally valid ways of conceiving the substance of the world, either involving substitutions for our categories, or their augmentation in, as many would feel, bizarre and unmotivated ways. What are the barriers to an entitlement to wood spirits, ectoplasm, gods and a plethora of existing but non-actual spatio-temporally unrelated concrete possible worlds?²²

That's a concern about whether entitlement can reach far enough to meet our needs. But there are also concerns about its potency even where it does reach. In general, it has to be recognised that the unified strategy can at most deliver a *sceptical solution*—so will disappoint those who are disappointed with sceptical solutions in general. Sceptical solutions concede the thrust of the sceptical arguments they respond to. Kripke's sceptical solution, for instance, concedes that meaning and its cognates are shown to be non-factual by the famous (putatively Wittgensteinian) sceptical argument. The unified strategy likewise concedes the basic point of the sceptical arguments to which it reacts, namely that we do indeed have no claim to know,²³ in any sense involving possession of evidence for their likely truth, that certain cornerstones of what we take to be procedures yielding knowledge and justified belief hold good. It then attempts to provide an accommodation with this concession, arguing that there is nevertheless no irrationality, or capriciousness, in our proceeding in the ways we do—that we are warranted in so proceeding but warranted in a different way. That is, of course, a very important claim if it is true. But there is no disguising the fact that the exercise comes as one of damage limitation. That will disappoint those who hanker after a demonstration that there was all along, actually, no real damage to limit—that the

22. The last is no mere rhetorical flourish. Lewis's ontology stands to (most of) ours as an ontology of other minds stands to solipsism.

23. Note: have no *claim* to know, not: do *not* know.

sceptical arguments involve *mistakes*. Good luck to all philosophers who quest for such a demonstration.²⁴

A more specific concern about potency is what I earlier termed the ‘leaching’ problem. 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?

The short answer is that there is leaching, but that it is at one remove and can be lived with. In general, to be entitled to trust that, for example, my eyes are right now functioning effectively enough in conditions broadly conducive to visual recognition of local situations and objects is to be entitled to claim that my vision is right now a source of reliable information about the local perceptible environment and is hence at the service of the gathering

24. Stephen Schiffer (see e.g. Schiffer [2003], pp. 68–9, and Schiffer [forthcoming]) usefully distinguishes between ‘happy face’ and ‘unhappy face’ solutions to paradoxes. Happy face solutions consist, broadly in the disclosure of mistakes in the premisses or reasoning of the paradox. But in the general run of philosophical paradoxes, Schiffer counsels us to ready for the possibility that there is no solution of this kind, that the only way out will be via conceptual revision. There is then a further division: *weak* unhappy face solutions propose conceptual revisions that are broadly conservative of the purposes and utility of the paradox-generating concepts; but in other cases—where only a *strong* unhappy face solution is possible—each of the possible paradox-preempting revisions will involve significant loss.

The general form of solution to the paradoxes of scepticism pursued by the unified strategy rather straddles the boundary between happy face and weak unhappy face. One can treat the sceptical arguments as involving a mistaken conflation of evidential justification and warrant—as overlooking the possibility of rational entitlement. Or one can see the invocation of entitlement as, in effect, a form of conceptual revision—extension—of our conception of the range of ways in which acceptance of a proposition can be justified. But in any case, the proposal does not wear the happiest of faces—that would only belong to a solution which somehow faulted the sceptical reasoning as applied to evidence.

of perceptual knowledge. To be entitled to trust that other humans have mental states whose character may be accurately discerned by applying our normal interpretative criteria to the things they say and do is to be entitled to trust that other minds can be known in standard ways. And to be entitled to trust in the soundness of a basic inferential apparatus—to anticipate a discussion of the status of fundamental rules of inference on which I have not here embarked²⁵—is to be entitled to regard its correct deployment as serving the generation of proofs and hence, since what is proved is known, to be entitled to claim knowledge of the products of reasoning in accordance with it. In general, the effect of conceding that we have mere entitlements for cornerstones is not uniformly to supplant evidential cognitive achievements—knowledge and justified belief—with mere entitlements right across the board but to qualify our claims to *higher order* cognitive achievement. I am right now in possession of a plethora of perceptual knowledge concerning occurrences around me. That is a claim which, if the unified strategy delivers as hoped, I will be rationally entitled to make. But in order to be able to *know* that it is true, I need (this is a closure step, of course) to be able to know the presuppositions of its truth, some of which—we are taking it—sceptical argument has put beyond evidence. So scepticism demands the surrender of higher order knowledge—the claim to know that we know. But entitlement, in the best case, promises to save the warrantability nevertheless of the first order claim to know. And maybe that is enough to be going on with.

Dissatisfaction may remain.²⁶ Let C be any cornerstone which sceptical argument persuades us is beyond evidence and let P be any ordinary, non-basic belief in the region of enquiry for which C is a cornerstone which, in line with the train of thought we just ran through, we are supposedly entitled to regard as knowledgeable nonetheless. Let warrant include both evidence and entitlement, and assume that, although closure across (known) logical consequence is qualified for evidential justification (as suggested at the end of Section I), it holds for warrant. Then the

25. For an initiation of discussion of the part entitlement may play in the epistemology of basic logic, see my [forthcoming].

26. The following train of thought is a version of an objection put to me by Sebastiano Moruzzi.

leaching problem, in a sharper formulation, is that the following trio of claims may all seem to be warranted:

First, if we run a risk in accepting C, then we run a risk in accepting P.

This seems merely to articulate an immediate implication of C's being a cornerstone for the class of beliefs typified by P.

Second, we do run a risk in accepting C.

—after all, entitled as we may be, the fact has not gone away that we have no evidence for C.

Third, P is known.

—a claim we are warranted (entitled) in making, by hypothesis. But now it appears that we must be warranted in claiming both that P is known and that we run a risk in accepting it. And that seems, near enough, a contradiction. A major part of the point of the concept of knowledge is that it is meant to mark a state in which belief is *safe*, in which it is risk-free. If it does not do that, what is the content of the claim to have knowledge that P?

The reply I am making in behalf of the unified strategy, transposed to this form of the leaching worry, is that what is wrong is not the third claim but the (consequent of the) first. What necessarily inherits the risk we run in trusting C without evidence is not our belief that P—for we may in fact have reliable evidence for P—but our belief that we have reliable evidence for it. To be sure, to *claim* to know P is indeed to promise that it is safe to accept P. However, that promise is not automatically worthless, or inappropriate, if the claim to know is not itself knowledgeable. It will still have every point if enough has been done to ensure that all that remains to put the knowledge claim at risk is the possible failure of conditions in which everybody, speaker and audience, (rationally) trusts.

One final matter. The discussion has proceeded with no mention of the opposition between *internalist* and *externalist* views of knowledge and justification. But its spirit, it may seem, has been very much internalist: entitlements, it appears, in contrast with any broadly externalist conception of warrant, are essentially recognisable by means of traditionally internalist resources—*a priori* reflection and self-knowledge—and are generally independent of the character of our actual cognitive

situation in the wider world²⁷—indeed, are designed to be so. Anyone who thinks that the paradoxes of scepticism are best solved, or dissolved, by proper emphasis on the external character of knowledge, or genuine warrant, is therefore likely to be impatient with the present project. If knowledge, and justification, are essentially *environmental*—are constituted by (perhaps reflectively inscrutable) contingencies of our cognitive powers and the way they enable us to interact with the external world—then no mere sceptical paradox, developed in the armchair, can show that we have no knowledgeable or justified beliefs. So why bother trying to make out entitlements?

Fully to address this reservation would need a complex and extensive discussion. But one immediate observation is that what is put in doubt by sceptical argument is—of course—not our *possession* of any knowledge or justified belief—not if knowledgeability, or justification, are conceived as constituted in aspects of the external situation in which we come to a belief. (How indeed could armchair ruminations show anything about that?) What is put in doubt is rather our right to *claim* knowledge and justified belief. It is this which the project of making out entitlements tries to address and which, on what seems to me to be a correct assumption, externalism is impotent to address.

That assumption is that epistemic values are subject to a division broadly similar to one within moral values. While some meta-ethical views—classical utilitarianism, for instance—can be seen as driven by a sort of moral monism, it is intuitively plausible is that there are at least two quite different kinds of virtue which an action may possess or lack: virtue of consequence (utility), and virtue of provenance—of conscience, or integrity—relating to the attitudinal states of the agent that determined her choice to act in that particular way. The two types of virtue are not, of course, independent—good conscience requires that one reckon with the foreseeable consequences of one's actions—and there seems no reason to expect that one should generally trump the other, still less to expect reducibility in either direction. I want to endorse a broadly analogous distinction in the ethics of belief: that we should allow a comparable kind of division between considerations of *intellectual integrity* and considerations to do

27. This claim would need qualification to allow for entitlements of substance.

with the *situational provenance* and other potentially fortunate or unfortunate aspects of the circumstances of a particular belief (for instance, its being the product of a reliable—truth-conducive—belief-forming mechanism). Both categories of virtue are important—indeed, I would argue, indispensable. So those philosophers who have done so have been right to lay stress on notions of knowledge, or justification, which emphasise the second. But, again, there is no reason to expect either type of virtue to reduce to, or trump, the other.

Descartes' project in the *Meditations* was one of harmonisation of his beliefs with the requirements of rational conscience and its timeless appeal is testimony to the deep entrenchment of virtues of intellectual integrity in our cognitive lives. The *right to claim* knowledge, as challenged by scepticism, is something to be understood in terms of—and to be settled by—canons of intellectual integrity. The paradoxes of scepticism are paradoxes for the attempt at a systematic respect of those canons. They cannot be addressed by a position which allows that in the end thoroughgoing intellectual integrity is unobtainable, that all we can hope for is fortunate cognitive situation. When good conscience fails, there are still, indeed, other good—circumstantial—qualities which our beliefs may have. But what is wanted is good conscience for the claim that this possibility is realised on the grand scale we customarily assume.²⁸

28. Versions of some of these ideas were presented in my NYU seminars on Scepticism in Spring 2002, and later in that year at a departmental colloquium at the University of Bristol, at the European Summer School in Analytical Philosophy held in Paris and at the Birkbeck Philosophy Society. They also featured prominently in a series of three seminars given at the University of Bologna in January 2004. My thanks to all who participated in those discussions, which generated innumerable improvements. I have also been greatly helped by the comments of my colleagues—Roy Cook, Philip Ebert, Nikolaj Jang Pedersen, Agustin Rayo, Marcus Rossberg, and Robbie Williams—in the Arché AHRB project on Foundations for Classical Mathematics, who have patiently allowed several of our weekly project seminars to be diverted onto this material, trusting in its eventual relevance to the issue of fundamental *a priori* knowledge. My thanks to Annalisa Coliva, Duncan Pritchard, Stephen Schiffer, and Tim Williamson for helpful comments on an earlier draft. I am especially grateful to David Enoch and Joshua Schechter who have been working independently [Enoch and Schechter, unpublished manuscript] on a generalised development of the broadly Reichenbachian direction, focused on the justification of belief-forming methods, and who each provided me with extensive constructive criticisms and comparisons with their own approach. I hope to take up the comparisons on another occasion. Most of the research for the paper was conducted during my tenure of a Leverhulme Research Professorship and I once again gratefully acknowledge the support of the Leverhulme Trust.

REFERENCES

- Burge, T., 1993, 'Content Preservation', *The Philosophical Review* 102, pp. 457–88.
- Enoch, D. and Schechter, J., 'How Are Basic Belief-Forming Methods Justified?' Unpublished TS.
- Evans, G., 1980, 'Things Without the Mind', in Zak van Straaten, ed., *Philosophical Subjects: Essays Presented to P. F. Strawson* (Oxford: Clarendon Press) pp. 76–116.
- Goodman, N., 1995, *Fact, Fiction, and Forecast* (Cambridge, Mass.: Harvard University Press).
- Hume, D. *An Enquiry Concerning Human Understanding* (ed. L. A. Selby-Bigge).
- McDowell, J., 1994, *Mind and World* (Cambridge, Mass.: Harvard University Press).
- Putnam, H., 1994, 'Sense, Nonsense and the Senses: An Enquiry Into the Powers of the Human Mind', *Journal of Philosophy* 91, pp. 445–517.
- Reichenbach, H., 1938, *Experience and Prediction* (Chicago: University of Chicago Press).
- Reichenbach, H., 1949, *The Theory of Probability*, 2nd edition tr. E. H. Hutten and M. Reichenbach, (Berkeley: University of California Press).
- Reichenbach, H., 1968, *The Rise of Scientific Philosophy*, Berkeley and Los Angeles: University of California Press).
- Schiffer, S., 2003, *The Things We Mean* (Oxford: Clarendon Press)
- 'Skepticism and the Vagaries of Justified Belief', *Philosophical Studies*, forthcoming.
- Strawson, P. F., 1959 *Individuals: An Essay in Descriptive Metaphysics* (London: Methuen).
- Sudbury, A. W., 1973 'Could There Exist a World Which Obeyed No Scientific Laws?' in *The British Journal for the Philosophy of Science* 24, pp. 39–40.
- van Fraassen, Bas. C., 1980, *The Scientific Image* (Oxford: Clarendon Press).
- Wittgenstein, L., 1969, *On Certainty* (Oxford: Basil Blackwell).
- Wright, C., 1985, 'Facts and Certainty', Henriette Hertz Philosophical Lecture for the British Academy, December 1985, in *Proceedings of the British Academy* LXXI, pp. 429–72.
- Wright, C., 2002, '(Anti)-Sceptics, Simple and Subtle: G. E. Moore and J. McDowell', *Philosophy and Phenomenological Research* LXV, pp. 330–48.
- Wright, C., 'Intuition, Entitlement and the Epistemology of Basic Logical Laws', *Dialectica*, forthcoming.