

11

On Epistemic Entitlement (II)

Welfare State Epistemology*

Crispin Wright

341. ... the *questions* that we raise and our *doubts* depend on the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn.

342. That is to say, it belongs to the logic of our scientific investigations that certain things are *in deed* not doubted.

343. But it isn't that the situation is like this: We just *can't* investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put.

(Wittgenstein, *On Certainty*)

In a paper written a decade ago,¹ I argued for the rational good-standing of a notion (perhaps realized in a number of distinct ways) of *non-evidential warrant* or *epistemic entitlement*, and for the possibility of using such a notion to fashion a unified response to two of the most traditional and disturbing forms of sceptical paradox.² In this paper, I further explore the potential significance of epistemic entitlements so understood, suggest some

* Precursors of this material were presented at a departmental colloquium at Brown in April 2009 and at the AHRC-funded *Basic Knowledge* project conference on Scepticism, held at the Arché centre in St Andrews in June 2009, when Patrick Greenough was the commentator. Thanks to the participants on these occasions for their comments. The *Basic Knowledge* project transferred to the Northern Institute at Aberdeen in September 2009 and focused on epistemic entitlement throughout the academic session 2009–10. Thanks to the members of the project for a huge amount of helpful discussion of the issues over many seminars. Special additional thanks to Carrie Jenkins, Luca Moretti, Aidan McGlynn, and Elia Zardini for invaluable written comments.

¹ Wright 2004a.

² What I respectively termed *Cartesian* paradoxes, which make play with some putatively undetectable scenario of systematic cognitive dislocation, and *Humean* paradoxes, which make a case for the

refinements of the earlier discussion, reassess the prototype of entitlement derivable from Hans Reichenbach's ideas about the justification of induction, and respond to a number of objections and difficulties for my proposals which have since surfaced in the literature.

By a 'non-evidential' warrant, I have in mind grounds, or reasons, to accept a proposition that consist neither in the possession of evidence for its truth, nor in the occurrence of any kind of cognitive achievement—for example, being in a perceptual state that represents it to one that P, or seeming to recollect that P—which would normally be regarded as apt to ground knowledge or justified belief that P. Still, a non-evidential warrant is warrant to accept a proposition as true—in some suitably qualified sense of 'accept'. In Wright (2004a), the proposed notion of acceptance was that of: *trust*. I shall continue to work with that proposal here.

11.1. Preliminaries: Hinges, Warrants, Claims to Warrant, Liberalism, Conservatism, Scepticism

Wittgenstein's metaphor of questions and enquiry as pivoting on 'hinges' raises the question, how to characterize the range of the propositions he seems to have had in mind—the 'hinge propositions'—and their role: how exactly does enquiry 'turn' on them? The examples in *On Certainty* are a mixed bag, but we can approach the kind of thing Wittgenstein seems to be gesturing at by focusing, first, on the notion of an *authenticity-condition*³ in the sense illustrated by the respective third propositions in these examples:

<i>Moore</i>	<i>Zebras</i>	<i>Red Wall</i>	<i>Stranger</i>	<i>Red Wall II</i>
1. Appearance of a hand in front of my face	1. Appearance of Zebra-like animals	1. Experience of a red-looking wall	1. S tells you that P	1. Experience of a red-looking wall
2. Here is a hand	2. Those animals are zebras	2. That wall is red	2. P	2. That wall is red
3. There is a material world	3. Those animals are not cleverly disguised mules	3. That wall is not a white wall bathed in red light	3. S is truthful	3. My visual system is working properly

unavoidability of some vicious form of epistemic circularity at the root of our procedures of enquiry. Examples of each of these genres of course vary in their detailed structure.

³ Cf. Wright 2007. In some other writings I have used the term 'presupposition' for this notion. But of course it is a major question whether, or in what sense, the satisfaction of conditions of the relevant kind is indeed in large measure *presupposed* in ordinary cognitive commerce. So the more neutral 'authenticity-condition' seems better. I think the notion, glossed as in the sequel, coincides with that of an 'anti-underminer' in Jim Pryor 2013.

In each case, the entry under 1 depicts one kind of basic evidence for the 2-proposition; and in each case the 3-proposition—the relevant authenticity-condition—is uncontroversially related to 1 and 2 like this: that doubt about it is rationally precluded on the part of one who proposes to take the evidence for the 2-proposition depicted by 1 as sufficient. Absent other relevant information, any doubt about the 3-proposition must tend, in a rational subject, to undermine the force of the evidence described in 1 for the 2-proposition. Such a doubt may, as in the case of *Moore*, *Zebra*, and *Red Wall*, involve doubt about the 2-proposition too. But in all cases, it must involve doubt about the significance of the evidence depicted in 1. One who doubts the 3-proposition could not rationally move to belief in the 2-proposition just on the basis of evidence 1.

We can characterize the illustrated idea more generally by invoking a very broad notion of a *cognitive project*, defined as a pair: a question, and a procedure one might competently execute in order to answer it. Thus there is a cognitive project associated with the question, ‘What’s the weather like today?’, which one can execute by looking outside; another cognitive project associated with the same question which one can execute in a windowless room by looking at the forecast in yesterday’s newspaper. There is a cognitive project associated with the question, ‘Will you spend Easter at home?’, which one can execute by an operation of whatever it is one does to know of one’s intentions. There is a cognitive project associated with the question, ‘Are there planets associated with that star?’, which one can execute by the operation of a radio telescope and a suitable interpretation of one’s findings. So, an authenticity-condition for a given cognitive project is any condition doubt about which would rationally require doubt about the efficacy of the proposed method of executing the project, or about the significance of its result, irrespective of what that result might be.⁴ That my visual system is working properly, or that yesterday’s weather forecast is likely to be accurate, are among the conditions for the authenticity of the two mentioned projects for determining the weather; that I am likely to have a normally lucid awareness of my intentions is an authenticity-condition of the project focused on my prospective whereabouts at Easter; and that my radio telescope is functioning properly, as well as a whole load of theory about electromagnetic radiation, are authenticity-conditions for the project about the distant star.

Authenticity-conditions thus include such things as: normal and proper functioning of relevant cognitive faculties, the reliability of instruments utilized, the amenability of the circumstances to the proposed method of investigation, the correctness of relevant theory, the soundness of relevant principles of inference utilized in developing and collating one’s results, the good standing of relevant concepts used in any aspect of the enquiry, and so on. Typically, one takes a broad sweep of such conditions

⁴ As the alert reader will have spotted, this qualification is crucial if we are to capture the intended notion. If it were waived, a prior opinion about the matter under investigation might rationally lead one to doubt the competence or significance of the project if that opinion happened to conflict with the upshot; and that would then have the effect that the very proposition which the project found in favour of would rank as an authenticity-condition of the project concerned.

for granted in cognitive projects ranging from the quotidian, like glancing at a clock to see if it is time to leave, to the more methodologically self-conscious, like a carefully controlled scientific experiment. For cognitive projects that involve defeasible evidence, as in the illustrated cases, absence of doubt about authenticity-conditions will be rationally required if the evidence adduced is to be credited with its intended bearing on the conclusions drawn. It is, for example, rational to take the look of a wall as good evidence for its colour only if one has no reason to doubt that the lighting conditions are such as to allow the real colours of surfaces to be revealed in how they look. Likewise, it is rational to take Stranger's testimony as good evidence of what is the case only if there is no reason to doubt that Stranger is inclined to tell the truth

Authenticity-conditions so characterized are specific to a given cognitive project. We get a better fit with some of Wittgenstein's characteristic examples in *On Certainty* if we generalize—if we consider authenticity-conditions that are common to a large sweep of cognitive projects of a given kind. That will ensure the feature that doubts about 'hinges' will have the much wider bearing that Wittgenstein adverts to when he speaks of their 'standing fast' for us as being 'part of our *method* of doubt and enquiry'.⁵ Such a doubt will be a doubt which ought—rationally—to ramify into a more general doubt of some sort: a doubt about *any* investigation that uses some relevant apparatus or relies upon on a certain kind of evidence, or a doubt about the good standing of *all* previous investigations of a certain kind, or about the very subject matter of a large class of investigations, or about the propriety of their methods.⁶

Some of these general authenticity-conditions⁷ articulate rationalizations of our most basic ampliative inference patterns. For example, the proposition,

'Nature is pretty uniform; there is a lot of natural law out there.'

acts as a hinge for ordinary inductive inference: doubt it and you cannot rationally move from a regularity manifest in a sample to its inductive generalization. The proposition,

'Others have mental states which are broadly manifest in what they say and do, and their circumstances.'

⁵ Wittgenstein 1969: §151—Wittgenstein's own emphasis.

⁶ We get a better fit with Wittgenstein's principal examples by generalizing in this way, but we do not get a perfect fit. He also cites cases like, arguably, 'I am called L.W.' where, he avers, '...there is no judgement I could be certain of if I started doubting that' (1969: §490). Another favourite example is 'I have never been on the moon' (as affirmed in 1950). Here too, he suggests, a doubt would have to ramify in unsurveyable and paralysing ways. But the reason why that is so—when it is so—is not because the propositions concerned play the role illustrated of authenticity-conditions in a wide class of enquiries but because error about such matters would rationally raise non-specific suspicions that large tracts of one's cognitive life hitherto may have been compromised. How could it have happened that I travelled to the moon yet have absolutely no recollection of it? And how could the technical possibility of such a thing have been kept secret from everybody—or do others perhaps know of it and keep it from me? And hasn't everyone always referred to me by that name?—Has that too been some kind of conspiracy, or am I misremembering that I have been consistently so-called?

⁷ What John Hawthorne 2005 terms 'heavyweight' propositions.

acts as a hinge for the ascription of mental states to others based on their behaviour, and physical and environmental circumstances: doubt it and you can rationally form no opinions about the mental states of anyone other than yourself. The proposition,

There is an external material world, broadly manifest in normal sensory experience.

acts as a hinge for ordinary empirical beliefs about the world around one; doubt it and you can rationally form no opinions about your local environment on the basis of apparent perceptual experience. And the proposition—the denial of Russell's Hypothesis⁸—that,

The world did not come into being five minutes ago, replete with apparent traces of a more ancient history, but genuinely has such a history, disclosed, in the round, in presently available data.

acts as a hinge for ordinary beliefs about the past: doubt it and you can rationally form no—or only precious few—beliefs based on apparent memories and other kinds of presumed traces.

What role do authenticity-conditions, and in particular general hinges, play in the *justificational architecture* of ordinary beliefs? What have come to be termed the respectively *liberal* and *conservative*⁹ answers to this question are best characterized with respect to quadruples consisting of a target proposition, P; some envisaged form of evidence or cognitive achievement, E; an authenticity-condition, A; and a species of epistemic warrant, W. The conservative view will be that the accomplishment by means of E of a warrant of type W with respect to P will require as an enabling condition that one be in W (or some other specified warrant state) with respect to A. The liberal will demur, holding that the attainment of W for P by means of E need rest on no positive form of warrant for A—that it will suffice merely if the agent has no warrant for not-A.

Suppose, for example, that W is knowledge, and the first three elements in the relevant quadruple are those in *Red Wall II* in the table above:

- (C) My experience is as of a red wall.
- (P) That wall is red.
- (A) My visual system is working properly.

Then liberalism with respect to this particular quadruple says that, in order to accomplish knowledge of the proposition P on the basis of the evidence, E, I do not need to know, or independently have adequate grounds for accepting A, that my visual system is working properly. It is enough merely that I have no antecedent reason to doubt it. Conservatism, by contrast, will hold that it is only in an epistemic context where I have some specified kind—the conservative will tell us what kind—of prior warrant

⁸ 'There is no logical impossibility in the hypothesis that the world sprang into being five minutes ago, exactly as it then was, with a population that "remembered" a wholly unreal past' (Russell 1921: 159).

⁹ I believe that Jim Pryor introduced this terminology.

for acceptance that my visual system is working properly that my evidence E has any tendency to support P, that that wall is red.

Obviously there is going to be space for such a clash of attitudes with respect to any quadruple of the relevant kind. Consider this version of the *Stranger* example:

- (E) I ask Smith the time and, glancing at his watch, he says, 'Eight o'clock'.
- (P) It's eight o'clock.
- (A) Smith is telling the truth on this occasion.

Let W again be knowledge. Then in order, the conservative will say, to achieve knowledge that P on the basis of Smith's testimony E, I must independently know, or have adequate grounds for accepting A, that Smith is telling the truth on this occasion. For the liberal, by contrast, knowledge that P can be achieved, in the best case (presumably, when Smith is as a matter of fact speaking knowledgeably), purely on the basis of Smith's testimony; no antecedent knowledge of or other form of warrant for Smith's truthfulness is needed.

The relativization of the issue to the appropriate kind of quadruples is important in this respect: it allows that even when E and P are fixed, the merits of a conservative, or liberal stance may vary as a function of A—the particular authenticity-condition in question—and W, the kind of epistemic warrant that is involved. One may wish to be conservative about some but not all authenticity-conditions for a particular cognitive project. And one may wish to take differing views about the requirements of different kinds of cognitive accomplishment. Knowledge, for example, may be taken to be more exigent than other forms of epistemic warrant exactly in that it demands warrant for particular authenticity-conditions in particular cases where other forms of epistemic warrant for the same proposition do not.

In any case, it is clear there can be no sensible *global* opposition between liberalism and conservatism. No one is going to be liberal right across the board, for all types of proposition, all kinds of defeasible evidence or cognitive accomplishment that may bear on them, all authenticity-conditions for that particular bearing, and all forms of epistemic warrant. Such a view would condone an open flood of epistemic irresponsibility. But equally, conservatism across the board threatens to set impossibly high standards for the acquisition of warrant. The interesting, disputed questions concern for which selections of the four parameters conservatism is appropriate, and for which selections the more relaxed stance of liberalism is perfectly rational, and why.¹⁰

At the time of writing, it is still true that most liberal-conservative debate has been focused on the nature of the warrant provided by a perceptual experience for beliefs about the local environment. What is striking is that the considerations which, respectively, most powerfully motivate the opposed views seem to belong at different

¹⁰ Perhaps, indeed, a fifth parameter is called for: that of the context of interests in which the evidence in question is assessed, or the relevant putative cognitive achievement takes place. Certainly, it is quite plausible enough that, under the aegis of a broadly conservative view of the role of a particular authenticity-condition, A, how much, and what quality, of independent evidence for A is required in order for belief in the target

levels. The most powerful consideration on behalf of *dogmatism*—liberalism about basic perceptual warrant—is that we do not wish to deny the title of warranted belief to opinions that children, and others who are relatively epistemologically innocent, form without considering, let alone marshalling evidence to discount the kind of possibility typified by tricky lighting or the artful disguise of mules. So it seems we think that the acquisition of basic perceptual knowledge, or other forms of perceptual warrant, doesn't require the kind of epistemological 'policing' of authenticity-conditions that conservatism demands. But it is different when we adopt a stance in which we undertake to *scrutinize* our claims to perceptual knowledge, or warrant. Then we seem obliged either to take a positive view of any authenticity-condition that may be entered into the conversation, or to qualify our claim to warrant. Standing before the zoo pen and inclined to claim knowledge that the animals before me are zebras, there is no ducking the question, 'So you are taking it that your visual system is functioning satisfactorily today, and that those animals have not been artfully disguised to prevent their identifiability by casual observation?' And the intuitively needed answer if the claim to warrant is to be sustained will be not that 'I have no view about those matters, and am not required to have one. I have been reading Jim Pryor. It's enough that I have no reason to doubt either condition', but rather, 'Of course'. Looking out of the window I say, 'Look! A Blue Jay has just landed on the bird feeder.' You, teasing no doubt, say, 'So you take it that you are capable of distinguishing a Blue Jay from a Robin?' Or 'So, you take it that your eyes are functioning properly today?' Or 'So you are taking it that there is an external material world?' To take a liberal view about any of these issues is to hold that it is not something on which one needs independent warrant for a view en route to acquiring perceptual warrant for the original claim about the Blue Jay. But when doubts, facetious or otherwise, are entered at the level of claims, it seems that they have to be addressed, rather than merely finessed, and I cannot rationally profess agnostic indifference about the questions raised. Rather, in order for me to go on thinking that I have perceptual warrant for the original claim, I need to be able rationally to discount such doubts. That they can be rightfully discounted is a commitment of my claim, even if not a necessary condition of my *having* the warrant concerned.

There seems, then, to be a puzzling but clear difference between what we want to say about the conditions governing the *acquisition* of perceptual warrant and what we want to say about the conditions under which such a warrant may *rationally be claimed*. And the conservatives seem to have the better of the issue when what is in question is the latter: you cannot rationally profess agnosticism about something you acknowledge to be an authenticity-condition for the acquisition of a warrant which you simultaneously enter a claim to have acquired.

proposition to be warranted may very well vary as a function of the interests of the believer (or the attributer, or an assessor, ...). It is not so clear, though, that such a change of context can transform a case where a liberal view is appropriate, that is, when *no* independent warrant is required for A, to one where conservatism is appropriate and independent warrant is required. But I shall not further consider this kind of complication here.

But what is this notion of a 'rational claim' to warrant—what kind of considerations underwrite such a claim if not just the having of the warrant? As a first approximation to an answer, I have in mind whatever one might relevantly enter into an attempt to substantiate the assertion, perhaps in the face of a challenge, that one is indeed warranted in accepting a certain proposition. On any broadly externalist conception of warrant, one may in fact be in possession of warrant for a certain belief and yet able, *in foro* as it were, to offer nothing by way of defence that one is. And notice that the same will go for second-order warrant—for warrant for the proposition that one has warrant for the acceptance that P. The second-order warrant too can be construed in such a way that possession of it need be nothing of which one can come to rational awareness just by reflection. This need not be to say that the notion of a rational claim to warrant, as here intended, is something that needs to be understood in *internalist* terms. In my view, 'internal' and 'external' are, in this context, somewhat blunt instruments, which aren't really suited to clarify very much. Claim to warrant is a second-order notion, no doubt, and what is claimed may no doubt be understood in internal or external terms. But the best way to get an intuitive handle on the notion that I intend is not to attempt to bring it directly under those categories, but to think of such claims as something to be assessed, and sustained or rejected, in a context of rational discussion and adduction of evidence, commonly recognized—very much as a claim to innocence, or guilt, may be discussed and assessed in the forum of a court of law. That of course is consistent with contextually variable standards of admissible evidence, and rules of debate.

More should doubtless be said to substantiate the notion of claim to warrant, but I will not undertake that here. I think it is clear that we have the notion—it is operative, for example, in the complaint one often hears about externalist responses to scepticism that while construing reliability of method used, or counterfactual sensitivity, for instance, as sufficient for the knowledgeability of a true belief heads off any possibility of sceptical demonstration, a priori, that knowledge is impossible, it leaves us no better placed to tell that we actually have it.¹¹ Moreover the most worrying sceptical challenges—or so I contend—are in any case in effect targeted at the level of claims to warrant: Descartes, for instance, is explicitly concerned with what part, if any, of his putative knowledge can be rationally claimed to be certain.

I have asserted that our intuitions tend to side with conservatism when what is at issue are which warrants we may rationally claim. Specifically, in claiming warrant for accepting the proposition that P based on a certain kind of ground or putative cognitive achievement, one will feel that one needs, if pressed, to be in a position to discount any doubt which, if sustained, would undermine that warrant. The liberal idea of a *mere* defeater—a proposition a doubt about which does indeed properly undermine, but to which one needs to take no form of positive epistemic attitude merely in order to *have* warrant of a certain kind, seems to have no place at the level of claims. But now

¹¹ This line of thought is elaborated in detail in Wright 2008.

if that is right, then liberalism has nothing to offer when it comes to addressing scepticism at the level of claims. Yet as we have noted, the great difficulty with conservatism, whether at first-level or the level of claims, is that—unless curtailed at some point—it threatens to set impossible standards for the acquisition of warrant, or the acquisition of the right to claim it, respectively. For if a subject is to be required to possess independent warrant for the satisfaction of a relevant authenticity-condition, or for the right to claim that it is satisfied, then that seems to demand the satisfactory completion of a prior independent cognitive project. And that in turn will have its own authenticity-conditions. If a conservative attitude is taken in turn towards them, then do we not launch a regress of which the upshot must inevitably be (first- or claim-level) scepticism?

That's a problem with those authenticity-conditions into which an independent investigation at least presents itself as a possibility. But it is not the full extent of the problem. The difficulty with the 'heavyweights', or cornerstones—the big, general hinges of the ilk of 'There is an external material world, etc.'—is that the very possibility of independent investigation is itself shrouded in sceptical doubt. That's the thrust of the 'I–II–III' or 'Humean' genre of sceptical paradox discussed in Wright (2004a). What cognitive project can I undertake in order to engage the question whether there is an external material world save one which assigns to my apparent perceptual experience the very evidential significance for which the existence of an external material world is an authenticity-condition? But, then, if, as conservatism may seem to require, I need to investigate that condition independently, I am stuck. What cognitive projects can I undertake in order to engage the question whether nature exhibits a wealth of natural law save ones which assign to my experience of natural regularities the very evidential significance for which the existence of natural law is an authenticity-condition? But, then, if, as conservatism may seem to require, I need to investigate that condition independently, I am stuck—just as Hume famously observed.

The promise of the notion of epistemic entitlement is that it offers to return a positive response to the question: How can it be epistemically rational to repose confidence in an authenticity-condition for whose satisfaction one has no evidence or other form of cognitive warrant? If such confidence *can* be rational, and if in particular it may be rationally placed in the cornerstones, then the strategy of response to scepticism opens up that I outlined in Wright (2004a). But what is worthy of additional note is that introducing the notion of entitlement into the dialectic also promises to defuse the tension between our apparent first-level liberal but claim-level conservative intuitions about the requirements of warrant. At least it does so if we can sustain a comparison between the benefits of entitlement and certain aspects of the notion of a *moral right*. Simply, an agent does not need to know her rights in order to have them. Indeed, she may have no conception of a right. And when she acts in ways that her rights mandate, her actions are in good standing even if she is unaware that they are so mandated or, though aware, unable to make out a cogent case that they are. Think accordingly of entitlements as determining *epistemic rights* and hence as having analogues of exactly

these features for those of our mental actions that consist in the formation and management of belief.¹² Then if, for instance, there is indeed a general entitlement to take it that, absent evidence to the contrary, one's sensory faculties are working normally in conditions broadly conducive to their effective operation, a young child, with no developed conception of sensory abnormality or illusion, who forms beliefs spontaneously in response to the promptings of her sense experience, is acting fully within her epistemic rights, and is, in that sense, justified in so doing, exactly as the dogmatist—liberal—first-level intuition requires. In other words, there need be no distinction, in terms of their predictions about what is required in order to acquire warrant in a particular case, between liberalism about the relevant mode of warrant acquisition and a form of conservatism that holds, *contra* the liberal, that independent warrant for relevant authenticity-conditions is indeed required, but that it is conferred by epistemic entitlement. Neither demands more of the epistemic agent than the other. On the other hand, this form of conservatism is in a position to expect our intuitive reaction to a question like, 'So you are taking it that your eyes are functioning normally today?'—viz. to acknowledge its epistemic (if not conversational) appropriateness and to feel that its truthful answer is 'Yes'—one is indeed so taking it and rationally committed to doing so—but also to have a sense that that answer is, normally, *in perfectly good order*. Can liberalism make out its right to a similar expectation?

Of course, our intuitive reaction is exactly what is embarrassed by scepticism. How can this be in perfectly good order?—How can one rationally take things on trust in the course of enquiry for which one has, and perhaps can get, no specific evidence? The entitlement project aims to explain how.

11.2. Other Notions of Entitlement

I was not the first philosopher to annex the terminology of 'epistemic entitlement' to a kind of warrant intended to contrast with the conscious possession of justificatory evidence. A number of leading philosophers, including Fred Dretske,¹³ Tyler Burge,¹⁴ and Christopher Peacocke,¹⁵ have all offered proposals that may seem, superficially at least, to be aiming the same terminology at the same target. Though it would take us too far afield to review the matter in detail, my own sense is that the differences run deeper than the similarities. I will illustrate by offering some quick points of comparison with Burge's notion, according to my understanding of it.¹⁶

¹² Evidential warrants, by contrast, seem more naturally conceived as conferring *epistemic obligations*: as determining what one *ought* to think.

¹³ Dretske 2000.

¹⁴ Burge 1993, 1996, and 2003.

¹⁵ Peacocke 2004.

¹⁶ Considerations of space here preclude comparisons with the notions of Dretske and Peacocke at even the following modest level of detail, though some of the points of contrast with the Burgean notion about to be observed apply to Peacockean and Dretskean entitlements too. A useful comparative overview of the four notions is given in Altschul 2011.

First, there are differences in the respective extensions of ‘entitlement’. Burgean entitlements are not, in the first instance, warrants to accept certain specified propositions as true but bestow a right, rather, to rely on certain belief-forming *capacities* or *sources*—perception, memory, testimony—in coming to beliefs. One might still say, to be sure, that a thinker is entitled to, for example, her perceptual beliefs in a sense derivative from a Burgean entitlement to rely on her perceptual faculties. But even then, it remains that according to the notion of entitlement proposed in Wright (2004a) and here, we are entitled not to our perceptual beliefs as such but to accept certain propositions which underwrite the authenticity of the processes which lead to them. The entitlements I am proposing operate among the authenticity-conditions of enquiry, not the outputs.

Second, and relatedly, while both notions contrast with justification, it is with differing notions of justification that they contrast. For Burge, perceptual beliefs, and others that are the products of capacities on which we are entitled to rely, lack justification. And this is because Burgean justification is restricted to independently specifiable reasons that the thinker could in principle articulate. I have no objection to constraining the notion of justification in this way. All the same, having no justification for a particular belief in this sense is still consistent with its being evidentially warranted, that is, with its being based on (non-propositional) evidence, or more generally with its being the product of specific cognitive achievements made possible by the faculties on which Burge holds we are entitled to rely. There is a perfectly intuitive, non-Burgean notion of justification whereby someone looking out of a window and forming beliefs about the scene outside on the basis of his experience is justified in those beliefs. Entitlement in the sense proposed in the present discussion contrasts even with this more liberal notion of justification: it is entitlement to take for granted, to trust, without evidence or relevant cognitive achievement of any kind. Burge’s preference is to restrict the idea of reasons for belief to cases where beliefs are formed/sustained inferentially. The Burgean notion of entitlement is enlisted in an attempt to explain how non-inferential but empirically based belief can be in rational good standing.¹⁷ It is thus, again, a notion that applies within the sphere of cognitive achievement. But the notion I am proposing is one that applies *in lieu* of cognitive achievement.

Finally although I am less clear than I would like to be about the details of the considerations that, in Burge’s view, ground our entitlements, it appears that they are hostage to a priori inscrutable contingencies (like there being no external world, or one’s being a lifelong brain in a vat), whereas our enjoyment of the species of entitlements which I am here canvassing is intended to be a matter that is determinable by philosophical reflection—which is as it must be if entitlement is to be used to respond to the seminar-room challenges of philosophical scepticism. (I will return to one challenge to the claim of the reflective accessibility of entitlements in Section 11.6 below.)

¹⁷ As are the Dretskean and Peacockean notions.

In any case, fashioning the materials for a response to scepticism is explicitly no part of Burge's project.¹⁸

The major point of similarity between the Burgean notion and mine remains, of course, the 'welfare state' aspect: for both of us, the epistemically entitled subject has no epistemic work to do to benefit from the entitlement—indeed, need have no inkling of his entitlement. It is his epistemic right to believe the propositions to which he is entitled, or which are generated by cognitive capacities on which he is entitled to rely, even if he has no conception of entitlement, or of his epistemic rights. In this respect, Burge and Wright coincide.¹⁹

11.3. The Reichenbachian Paradigm Reappraised: The Problem of Attitude

Let me turn now to review an objection to what is, in effect, the prototype of the notion of entitlement—he did not of course use this terminology—developed in Reichenbach's attempt to develop a decision-theoretic response to inductive scepticism. We can begin once again with this toy example:

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 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 nutritious, he survives by eating it and will not otherwise do so; if the fruit is non-nutritious, 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.²⁰

Here is a simple summary of Crusoe's options and their respective foreseeable outcomes:

	<i>NUTRITIOUS</i>	<i>NON-NUTRITIOUS</i>
<i>EAT</i>	SURVIVE	DIE
<i>DON'T EAT</i>	DIE	DIE

¹⁸ Dretske 2000: n. 13 is similarly unconcerned with scepticism. Peacocke, by contrast, marshals the resources provided by his notion into a complex theoretical case that it is rational to discount sceptical hypotheses.

¹⁹ Indeed, this is a point of commonality with the Dretskean and Peacockean notions too.

²⁰ From Wright 2004a: 178–9.

Eating the unfamiliar fruits dominates—does better than all alternatives in one possible scenario and no worse than them in all others. So EAT is rational.

Reichenbach's master thought was that a practice of inductive inference might be justified along broadly similar lines:

	NATURE IS UNIFORM	NATURE IS HAPHAZARD
INDUCTIVE PRACTICE	MANY TRUE AND USEFUL BELIEFS	FEW TRUE AND USEFUL BELIEFS
OTHER PRACTICE	FEW TRUE AND USEFUL BELIEFS	FEW TRUE AND USEFUL BELIEFS

Inductive practice dominates—does better than all alternatives in one possible scenario and no worse than them in all others. So INDUCTIVE PRACTICE is rational.²¹

Reichenbach's own central concern was with the justification of inductive inference to probabilities, conceived as limiting frequencies of specified event-types in open-ended populations of events. As is familiar (and as Reichenbach himself was vividly aware), there are great difficulties in vindicating specific and natural-seeming rules of probability-projection just on the basis of the master decision-theoretic thought.²² In Wright (2004a), however, I was concerned with a much more general

²¹ Although the question does not matter for our present purposes, it is open to interpretation whether Reichenbach was concerned to argue that induction is a dominant strategy. Often his thought seems to equivocate between that claim and the idea that if any strategy will work, induction will. But here are three examples he offered to illustrate his intuitive point which suggest that dominance was the intended point—that if any method will work, it can *only* be induction:

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: 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: 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: 349)

²² For details about this problem for Reichenbach, and an attempt to resolve it on Reichenbach's behalf, see Wesley C. Salmon (1991).

worry—a worry about the benefits that might rationally be achievable, even in the best case, by reliance upon a pattern of inference that is underwritten only in a broadly Reichenbachian way. It seemed to me that there was a damaging problem concerning the kind of *attitudes* that a reliance, so rationalized, could underwrite—a problem affecting both any relevant cornerstones and the propositions issued by the inferential practice supposedly legitimized.

The problem can be brought out by a simple example. Suppose you need to measure the dimensions of a table. You have a tape measure, but for some reason a doubt has been entered about the accuracy of its calibration, with the result that you are now agnostic about whether it is suitable for the purpose in hand. Then clearly it would be irrational to reason like this: ‘Well, the tape measure is all I’ve got—if it’s accurate, or accurate enough, its results will be reliable; and if it’s not, well I have no other way of obtaining reliable results. So I should just use the tape measure and accept its results.’ At least, that is irrational if, by ‘accept’ the results of measuring with the tape measure, you mean: *believe* them—consider them to have been ascertained. You may still rationally accept them in a conditional spirit, or as a working assumption. But what you may not rationally do is to combine the attitudes of confidence in the results of the measuring with open-mindedness about the accuracy of the tape measure.

The moral is that even if it were possible to show that dominance reasoning, or some other form of decision-theoretic justification, does provide a kind of instrumental justification for the use of inductive methods, the most that would be delivered thereby would be an acceptance of the products of those methods *in the same spirit*, whatever exactly that spirit is. So if the desired output is *rational confidence* in the products of some belief-forming method, a Reichenbachian justification of that method is not enough. If rational confidence is to be the end product, then some form of rational confidence has to be incorporated into one’s acceptance of the methods. The attitude to the local hinges and cornerstones has to be one of *non-provisional confidence* if belief in the products is to be rationalized in turn. At the same time, it needs to be a rational attitude to take for reasons other than our possession of evidential support for the effectiveness of the methods concerned.

This was the train of thought that led me to propose that any useful form of entitlement had to license *rational trust*; and since it seemed that any rationale provided by Reichenbach for proceeding inductively would at best be consistent with utter open-mindedness about the likelihood that doing so would give good results, I discarded the Reichenbachian paradigm and moved to consider other possible grounds of entitlement. But I now think that rejection was premature. The Reichenbachian paradigm can handle the problem of attitude after all, and indeed quite simply. All that is required is that a state of trust be appropriately written into the decision-theoretic matrices. For example, the matrix for basic inductive inference may be refashioned like this:

	NATURE IS UNIFORM	NATURE IS HAPHAZARD
TRUST IN THE TRUTH- CONDUCTIVENESS OF INDUCTION	MANY TRUE AND USEFUL BELIEFS	FEW TRUE AND USEFUL BELIEFS
LACK OF TRUST	FEW TRUE AND USEFUL BELIEFS (OR MANY TRUE AND USEFUL BELIEFS, BUT AT THE COST OF THE RATIONAL INCOHERENCE OF COMBINING THEM WITH LACK OF TRUST IN THE METHODS WHEREBY THEY ARE ACQUIRED)	FEW TRUE AND USEFUL BELIEFS

And now—always provided we can justify the entries into the four boxes—we have a vindication of the rationality of trusting that induction is reliable and hence, absent conflicting evidence, of believing the generalisations and predictions it affords.

There is a natural immediate concern with this way with the problem. It is that justification following the Reichenbachian paradigm properly bears only on *voluntary action*. The paradigm argues that it is rational to perform a certain action, or to adopt a certain strategy for action, on the grounds, broadly, of expected utility, of one sort or another. The stage setting is one of determining what one should *choose* to do, in circumstances of uncertainty. But while, up to a point, and *pace* Hume, we can choose to infer in accordance with certain canons of inference rather than others, voluntarism about doxastic attitudes is a controversial, and widely disbelieved view. We do not, and cannot, simply decide to believe things or, weighing the evidence, choose to believe one rather than another proposition in the light of it, and the point is only the more plausible when the attitude concerned is trust. To be sure, I may in certain circumstances ask that you ‘trust me’, and you may respond by, as we say, putting your trust in me. But that means: writing the assumption that I will in fact perform in certain ways sensitive to your interests into the considerations that enter into your own decision-making. It does not mean: voluntarily falling into an attitude of calm expectation that I will in fact so perform. Or if it does, that is not something that you can just decide to do.

However, I do not think that the revamped Reichenbachian argument-schema is hostage to this concern. The question, remember, is whether our actual trust in induction, and other targets of sceptical doubt, is rational—not whether it was rationally entered into. If the notion of a thinker who could mould his doxastic attitudes at will is at least conceptually coherent, the standards for the rationality of the attitudes that he freely chose would presumably be unaffected, and he could properly be encouraged to

choose to form beliefs, or trustings, only when there was a sufficient rational case for doing so. We cannot do that, but it is still appropriate to assess the attitudes in which we find ourselves entrenched by the very same standards of rationality that we would apply to those of our imaginary doxastic free agent. If there are considerations, of a broadly decision-theoretic sort, that show that it would be rational to *choose* to trust in certain things, and form beliefs accordingly, if one could but do so, then presumably it is also disclosed as rational if one happens to be disposed to have the very same trustings, and propensities of belief-formation, as a matter of natural inclination, or training, rather than choice.

11.4. Leaching and Alchemy

That addresses the objection to the Reichenbachian strategy sustained in Wright (2004a). But the issue still arises, of course, about how, if all I can repose in inductive methodology, or in the cornerstones that underwrite it, is a kind of evidentially unsupported even if rational trust, I can nevertheless end up with something *epistemically superior* to evidentially unsupported trust when it comes to the outputs—the beliefs that inductive methodology, or operation within the evidential parameters set by the relevant cornerstones, allows me to form. That is what, in Wright (2004a), I called the ‘leaching’ problem. If all we have at the foundations of our belief-forming practices is trust, even if rational trust, how can elements in the superstructure acquire an epistemic status superior to that?

Wright (2004a) focused on the following way of developing this concern.²³ Let C be a cornerstone for a region of belief—say, ‘There is an external material world’—whose only epistemic credentials are that it may be rationally trusted. And let P be any quotidian proposition in that region of belief—say, ‘Here is one hand’, affirmed on the basis of sense experience in what are presumed to be good conditions of observation and observer—which we should ordinarily regard as known. Since C does not allow of knowledge, or evidence, there has to be a sense in which one runs an epistemic risk in accepting C, even if one trusts it rationally. In general, whenever the best one can do is place trust in something, however rational it may be to do so, there will be an element of precariousness which knowledge, or strong evidence, would forestall or attenuate. But then it may seem that, as is implicit in the very metaphor of a cornerstone, this precariousness must be inherited by the superstructure, so that we have, for any ordinary proposition P in the region in question which we should normally regard as known, the following conditional:

If there is epistemic risk in accepting C, then there is epistemic risk in accepting P.

Given, then, that there *is* epistemic risk in accepting C, and given that that kind of epistemic risk is exactly what is foreclosed by knowledge, what space is left for knowledge of P?

²³ Due to Sebastiano Moruzzi, in seminar discussion. Cf. Wright 2004a: 208–9.

My reply to this concern was that the conditional is wrong. Since the risk involved is that of acceptance of a proposition without knowledge of or evidential warrant for its truth, the idea that the risk transfers from C to P is just the contrapositive of the idea that knowledge and evidential warrant transmit from P to C. So interpreted, then, the worry about leaching is just a version of the discredited assumption that warrant is unrestrictedly transmissive.

As I observed, however, two qualifications of, or observations about, this reply seem to be needed. The first is that a leaching phenomenon has to be acknowledged at higher order. If the full extent of the epistemic credentials of the proposition that there is an external material world is that it is something in which one may rationally place trust, then that is also the full extent of the epistemic credentials of the proposition that sense experience provides a quotidian sort of knowledge about an external material world. So our *claim* to perceptual knowledge is going to be a 'mere' entitlement even if—if we do indeed have it—the knowledge itself is, naturally, something more. And so for putatively knowledge-acquisitive methods in general. If, for example, as I have argued elsewhere,²⁴ our belief in the validity of our most basic rules of inference is likewise only a matter of (mere) entitlement, then while such rules may indeed be at the service of extending our knowledge—when relevant premises are known—we will at best be entitled to claim to know that they are so, rather than knowing that we know, and hence would have no second-order knowledge of the conclusions to which they lead in any particular case.

The second qualification is to acknowledge that once it is allowed that it is possible to know or possess evidence for a proposition for certain of whose consequences one has only an entitlement, we *are* implicitly jettisoning certain forms of closure principle. Closure for warrant *in general* is unimpugned, provided warrant in general embraces both entitlements and evidence. But closure for *specific types* of warrant is discarded. That is, it is denied that whenever one possesses a certain kind of warrant for a proposition, that ensures that very same kind of warrant is available for anything entailed by that proposition. But this discard may seem well motivated in any case. The appearance of the animals in the cage is, in normal circumstances, evidence that they are zebras, but not, plausibly, that they are not mules so artfully disguised as to be observationally indistinguishable from zebras. Surely, observation alone cannot warrant the thought that one, rather than the other, of a pair of observationally indistinguishable situations obtains. Likewise, if you tell me that P, I can have a testimonial warrant to accept that P is true. But do I thereby have even a partly testimonial warrant for thinking that you told me the truth? The warrant provided by a piece of testimony does not, surely, extend to the reliability of that very piece of testimony.

Now, these ideas come under some pressure from a kind of dual of the concern about leaching: a concern about what Martin Davies nicely termed a kind of epistemic

²⁴ Wright 2004b.

alchemy.²⁵ Just as one might worry that a mere entitlement might, as it were, leach up from the foundations of our system of belief into the province of what we want to regard as genuine cognitive achievement, so one might wonder how genuine cognitive achievement might be prevented from driving down into the sphere of entitlement, so that a proposition, C, our entitlement to which underwrites our evidence for a more specific proposition P, might stand to then inherit, as it were, P's evidential credentials should there be an entailment back from P to C—with the lead of an original entitlement thereby mysteriously transformed into the gold of genuine evidential warrant. This would seem to be a kind of entitlement-conservative counterpart of the 'easy knowledge' problem that afflicts the liberal.

To be sure, the immediately available reply to this concern is exactly the same as to the concern about leaching, viz. that leaching and alchemy require specific forms of closure principle which the theorist of entitlement should reject—indeed, which anyone should reject who acknowledges the phenomenon of warrant transmission failure, one might think (and is savvy enough to distinguish that matter from closure for undifferentiated warrant).

The question, though, is how comfortable we can be with the resulting position. A reminder is worthwhile of the way that Davies originally elaborated the concern:

Ordinarily, we think that, if I review some of my beliefs, $P_1; \dots; P_n$, and notice a valid argument from those premises to Q then I should adopt the belief Q or, if other considerations argue against Q, then I should reconsider my beliefs $P_1; \dots; P_n$. If there are warrants for me to believe $P_1; \dots; P_n$ then, if I also believe Q, I shall again believe something for which there is a warrant. I shall think the thing that is the thing to think. But there is a distinction between believing something that is, as it happens, the thing to think and believing something because it is the thing to think. If I believe $P_1; \dots; P_n$ because there are warrants for doing so, then I do well doxastically. If I start out believing $P_1; \dots; P_n$ because there are warrants for doing so, and I go on to believe Q precisely because it follows from those premises, once again I do well doxastically.

These familiar thoughts suggest that, given the obvious entailment in the I–II–III argument, if we believe the type-II proposition that is supported by the evidence described in the type-I proposition, then we should also believe the type-III proposition that is the argument's conclusion. If considerations about non-transmission argue for going no further than the antecedent trust in the type-III proposition then we should reconsider whether belief is the proper attitude towards the type-II proposition.²⁶

A possible reply would be to say that the ordinary, 'familiar' thoughts about these matters to which Davies adverts make nothing of the theoretical distinction, recommended in Wright (2004a), between acceptance in general and belief in particular; and that the ordinary idea that one should believe the known consequences of what one believes—or revise one's beliefs accordingly—really has no axe to grind against the revised suggestion that one should accept responsibility for the consequences of one's

²⁵ Davies 2004: 222.

²⁶ Davies 2004: 222–3.

beliefs, or acceptances in general, by accepting them too, in whatever mode of acceptance is appropriate to them—or again, revise one's beliefs accordingly.

That, though, is to miss the point. The question is, what account can be given of the *rationality* of the acceptance of the consequences of one's beliefs once one cannot simply say that they will inherit whatever epistemic credentials belong to those beliefs—that, as consequences, they are mandated by exactly the same considerations?

It's a good question. In response, one might stonewall: one might offer that it is simply a primitive facet of responsible doxastic management to accept the consequences of what one accepts: that acceptances are rationally controlled by in-rules of warrant and out-rules of commitment, as it were. But perhaps it is possible to say something a little more illuminating than that. For closure of evidential justification is going to fail only in cases where transmission fails. If it is right that cases of transmission failure are one and all cases where some kind of independent good-standing for the conclusion is required to underwrite the presumed warrant for the premises, then a rational thinker would be committed to taking that view of the conclusion just by crediting herself with rational justification for the premises, independently of the entailment. So she ought to acknowledge the conclusion anyway.

If that is a satisfactory reply, it meets the concern as expressed by Davies: we can explain why we do 'doxastically well' in accepting the consequences of evidentially justified beliefs even when the particular type of justification concerned does not transmit. However, in his contribution to the present volume²⁷ Aidan McGlynn rehearses an argument to prove that the resulting position is unstable in any case. McGlynn reminds us that there is at least one restricted, very intuitive form of closure principle for evidential warrant which will suffice to ensure alchemical transformations—or conversely, leaching—in cases where the entitlement theorist, as so far characterized, should wish to resist it: cases where a firewall has been proposed between the domain of entitlement and the realm of evidence. The principle in question is what McGlynn calls:

ClosOR: If one is (evidentially) justified in believing *P* and one knows that *P* entails *P* or *Q*,²⁸ then one is justified in believing *P* or *Q*.

This principle is seemingly *per se* alchemy-inducive. At least it is so, provided it is accepted that evidential justification is closed across a priori known equivalence—*Equivalence Closure*. Here is the argument:

Alchemy-conjuring: Suppose *P* entails *Q*. Then *Q* is an a priori equivalent of *P* or *Q*, and it will follow by Equivalence Closure that there will be evidential justification for *Q* whenever there is evidential justification for the disjunction. *ClosOr* will then ensure that there is evidential justification for the disjunction whenever there is evidential justification for *P*. Thus, letting *P* = 'Here is a hand' and *Q* = 'There is an external material world', it will follow that if there is evidential justification for the former, there is *evidential* justification for the latter.²⁹

²⁷ See Ch. 9.

²⁸ Disjunction is here understood inclusively.

²⁹ This argument is given by John Hawthorne, at pp. 39–40 of Hawthorne 2004.

And, of course, the same goes for *Zebras* and all the other usual suspects. Thus let $P =$ 'Those animals are zebras' and $Q =$ 'Those animals have not been cleverly disguised to look just like zebras.' Then since P or Q is a priori equivalent to Not (not- P & not- Q), the warrant for P or Q furnished by ClosOR will ensure, by Equivalence Closure, that there is evidential warrant for: 'It is not the case that those animals are not zebras and have been cleverly disguised to look just like zebras'³⁰—which contains no reference to mules, to be sure, but is just as good as Dretske's original conclusion for eliciting the intuition of evidential closure failure. For it impresses as ungainsayable that, just by looking at the animals, I cannot get evidence that they have not been disguised in ways undetectable just by looking.

This impresses as a paradox. For the two principles deployed—ClosOR and Equivalence Closure—may also seem ungainsayable. As McGlynn rightly emphasizes, ClosOR seems compelling not because one thinks that, given evidential justification for P , one is sure to have *some kind* of evidential justification for P or Q , but because one thinks that one is bound to have *the same* justification for it. Surely, justification for either of its disjuncts *constitutes* justification for a disjunction. Surely, wherever and for whatever reason failures of evidential closure can occur, inferences across a single disjunction introduction step will not be one such case.

What about Equivalence Closure? So far as I know, up until very recently no one had ever unblinkingly suggested that evidential warrant may fail of closure—or transmission—across a priori known equivalence. But in the presence of paradox, perhaps it is worth asking why not. And actually, if we allow that equivalent statements may nevertheless embody differing units of information, and if evidential relations are held to be sensitive to such informational differences—if they are *hyperintensional*—there is no immediately evident reason why not: whatever account is proffered to explain failure of evidential closure, or more specifically transmission, across entailments from stronger to weaker contents might be expected to be generalizable in principle to differing, even if equivalent, contents. Such a proposal is worked out in interesting detail in Wesley Holliday (forthcoming).³¹ But, although this development is highly significant if it can be sustained, it is not the direction I shall take here.³²

³⁰ Indeed, no need to concern ourselves with the detail—e.g. the logical form—of Q . When P entails Q , Q is equivalent to P or Q irrespective of its detail. So if ClosOR and Equivalence Closure are good, evidential closure (and hence in general Transmission) has to be good too.

³¹ The appendix to Holliday's paper applies his proposal to examples of the kind presently concerning us. Holliday's ideas draw on but differ in detail from recent work of Stephen Yablo. See Yablo (forthcoming).

³² In fact, we can probably finesse the issues about Equivalence Closure in the present context. McGlynn has an argument to the same effect as *Alchemy-conjuring* that avoids reliance on Equivalence Closure, and appears effective enough. His actual formulation (Ch. 9, this volume) is hostage to issues about the proper characterization of cornerstone propositions, but it seems to me that this is avoidable, and that the essential gist is as follows. Let P entail Q . Suppose Q is a cornerstone for the acquisition of a certain kind of evidential justification, e, for P , and that you have as yet no justification for P , but that you do have an entitlement to Q . Given closure for warrant in general—which I accept—you must already have *some kind* of warrant for P or Q since it is entailed by Q and you have an entitlement for Q . But these assumptions seem insufficient to guarantee that you have any *evidential justification* for P or Q . So it must be that you have an

My own inclination, as McGlynn reports, is to admit a kind of alchemy after all.³³ That is, I propose to allow that the (visual) warrants in the classic examples do, after all, transmit. This is not, however, to surrender the conservative instincts that fuelled the diagnosis of transmission failure in the first place, still less to lapse into a Moorean liberalism. It is useful here to distinguish between two notions of a warrant-transmissive argument.³⁴ There are arguments that, relative to a certain ground for the premises, and certain kinds of collateral information,³⁵ are apt to confer a *first-time* warrant for acceptance of their conclusions; and there are arguments that, relative to a certain ground for the premises, and certain kinds of collateral information, are apt to *enhance an anterior* warrant for accepting their conclusions. Admitting alchemy in the case of *Moore* and others of the usual suspects need involve no compromise of their diagnosis as exhibiting failure of transmission of warrant of the first kind: as being useless for the purposes of surmounting an antecedent lack of warrant for their conclusions. Since, or so I believe, liberalism about the conclusions in question *qua* authenticity-conditions implies that first-time warrant *can* be conferred in these cases, this diagnosis continues to require that liberalism is false.

Still, the possibility is left open that the arguments in question can serve to enhance an anterior warrant for their conclusions. The admission of alchemy is the admission that they do. But the question is, what kind of enhancement? We are now allowing both that some kind of antecedent warrant to discount funny business at the zoo is prerequisite for rationally taking the appearance of the animals at face value—prerequisite for the evidential significance we assign to the data when we take the appearances to support the claim that the animals in the cage are zebras—and, by ClosOR, that the significance so assigned may then be enlisted in support of the disjunction, ‘Those animals are zebras *or* those animals have not been cleverly disguised to look just like zebras’—and hence in support of any equivalent of it. So *new evidence* is acquired for ‘It is not the case that those animals are not zebras and have been cleverly disguised to look just like zebras’. But now: is the *rational credibility* of the latter enhanced by this run-around? Should one be more confident in it than before?

The conservative should say ‘No’: that whatever limit there was to the rational credibility of that proposition as an object of trust, say, or as an object of belief based on independent evidence, remains as an upper bound on its rational credibility as an object of belief based on the appearance of the animals in the cage. To suppose that its credibility could somehow be raised higher than that would be to suppose that the visual appearances could somehow give more support to a proposition than there is independent reason to take them to be reliable or significant. That would be truly alchemical, and seems absurd. What, it may be suggested, *is* possible is rather that one’s

entitlement to it. However, once you exploit this entitlement to acquire evidence *e* for *P*, this evidence, it seems must become *evidence* that *P* or *Q*, by ClosOR. That’s an alchemical transformation.

³³ Ch. 9, pp. 184–7.

³⁴ Compare Wright 2012: 451–83.

³⁵ See Wright 2012 for more on the need for this proviso.

epistemic situation changes, but *without improvement*. In the abstract, there is nothing *outré* about the idea that one may add to one's evidence for a certain proposition without making it any more credible—without improving one's epistemic situation with respect to it. This is already perfectly intuitive in cases where one already possesses strong evidence for a certain belief, and then adds to it a measure of evidence of lower grade—say, one has just accomplished a proof of some theorem and then gets the testimony of a well-meaning but less able logician that he has proved it too. It is perhaps more startling if one can accomplish evidence for the first time for a proposition without improving one's antecedent epistemic situation with respect to it—if one may acquire evidence for a proposition for which antecedently one merely had reason to trust without thereby enhancing the rational credibility of that proposition. Well, maybe that possibility is exactly what we need to recognize.

There is a natural but confused objection to this line. Consider the situation of the disjunction, P or Q , before one gets any visual or other evidence that P ; and suppose that at that point, one's credence in P is 0.5. Let Q be an authenticity-condition for the project of looking to see whether P , and suppose, supported by an entitlement, that one's credence in Q is high—say 0.95. Then, when one does get evidence that P , supported by Q , one's credence in P will naturally rise. Won't that force a rise in the credence one should place in the disjunction? Not so. This line of thought forgets that Q is entailed by P , so that the probability of P or Q is just that of Q , and hence will not rise, even if that of P does, unless the latter goes higher than the prior probability of Q . But the conservative point is exactly that that cannot happen.

As McGlynn acknowledges, it is an open question whether the concession of alchemy involves in the end any fundamental awkwardness for conservatism. Clearly, if we take the line I have just sketched, we must be careful in characterizing what it is for an argument to involve a failure of transmission of warrant in the second, more generic sense distinguished above. When transmission fails in that sense, the key feature has to be not that a given evidential warrant for the premises of a valid argument simply doesn't constitute an evidential warrant at all for its conclusion, but rather that the significance rationally assigned to the evidence concerned is bounded by one's anterior confidence in the conclusion, so that the argumentative routine in question is not available to enhance that prior confidence. I am, of course, aware that this is a significant reconception of the way that I have tended to characterise transmission failure in some previous work, but it would take me too far afield to explore its implications here.

'But what of the plausible thought that one simply doesn't have any *visual* warrant for taking it that the animals concerned are not cleverly disguised mules?' I think the reflections above force us to say that the plausible thought involves a conflation. That is: we don't have a visual warrant for thinking that *those animals have not been cleverly disguised in a visually undetectable way*, but we do, in the relevant circumstance, have a visual warrant for thinking that *those animals are not mules that have been so disguised*. Maybe we are confused by the operation of some kind of

implicature here: maybe saying, or thinking, ‘It is not the case that those animals are cleverly disguised mules’ somehow implicates, in any context of a certain (normal) kind, that ‘Those animals have not been cleverly disguised’. But anyway, it doesn’t entail it: not-(P&Q), dear reader, does not entail not-Q! The sense remains strong to be sure, that if I am about to take the visual evidence for the presence of zebras at face value, I should be given pause if someone brings up the cleverly disguised mule possibility—that that will need an independent answer, beyond the already available visual evidence. But that point—that there will be a need for independent assurance—is accommodated by the acknowledgement that the strength of the visual evidence against the disguised mule possibility is bounded by the prior credibility of the proposition that it does not obtain.

If alchemy, so qualified, is allowed, then there is a corollary—albeit a dialectically inconsequential one—for the discussion of the leaching problem. The threat of leaching depended on evidential closure plus the assumption that there can be no evidential warrant for cornerstones. My response in Wright (2004a), outlined earlier in this section, was to reject evidential closure. But now we have in effect admitted evidential closure after all. So has the leaching problem come back too? Well no, since we no longer have the assumption in place that there can be no evidential warrant for cornerstones! What remains true, though, is that such evidence as we can claim for cornerstones depends for its force on our antecedent reason to trust in the truth of those very same propositions, so that the *ultimate* authority we have for accepting them depends on the rationality of that trust. That was, in effect, the concession I made about second-order leaching, and it still needs to be made. The admission of alchemy allows that evidential warrant is closed. But our warrant to take it that our basic evidential warrants have the significance we characteristically attach to them remains a matter of entitlement.

11.5. Pragmatic vs. Evidential Warrant

Among the first critics of the entitlement-strategic way with scepticism proposed in Wright (2004a) was Duncan Pritchard, who wrote as follows:

...this is, at best, merely a pragmatic resolution of the sceptical paradox, since it simply notes that accepting the sceptical conclusion would lead us to absurdity and intellectual stasis and then argues on this basis that we must reject it and therefore accept the legitimacy of our ultimately groundless believing. That not believing in hinges would be intellectually self-subverting in this way is not, however, an epistemic reason for thinking that such beliefs are true, but merely a pragmatic consideration which counts in favour of our proceeding as if they are true. Non-scepticism is thus defended on the grounds that it is the practical alternative, but we knew that already. Despite the well-advertised claims to contrary, then, this particular hinge proposition thesis presents us with no *epistemic* response to the sceptic at all.³⁶

³⁶ Pritchard 2007: 207 (my emphasis).

Carrie Jenkins sympathizes:

It does indeed seem that, if all Wright is doing is defending the practical rationality of accepting S, then his project does not promise any kind of solution—sceptical or otherwise—to Cartesian scepticism. For Cartesian sceptical argument appears to target the thought that trust in propositions like S is epistemically rational, not the thought that it is practically rational. Think of the ways Cartesian sceptical doubts are motivated: we are encouraged to consider, for instance, that an evil demon might be manipulating our sensory input so as to deceive us. How does this affect our views as to the practical rationality of trusting our senses? We might well think that it makes no difference at all: that, as far as we can tell, what we have most practical reason to do is surely to trust our senses and hope that we end up with true beliefs (since if they were reliable we'd have most practical reason to trust them, and we've no idea what we'd have most practical reason to do if they were in fact radically unreliable). The evil demon thought experiment is interesting because it is supposed to make us reassess the *epistemic* rationality of trusting our senses. In short, then, if all Wright is showing is that we are 'empowered' to 'dismiss' cognitive dislocation scenarios in the sense that, as far as we can tell, we have sufficient practical reason for so doing, nothing has been done to 'nip Cartesian sceptical argument in the bud'. For Cartesian sceptical argument does not engender concerns as to whether it is practically irrational to dismiss these scenarios, only concerns as to whether it is epistemically irrational.³⁷

Pritchard (and Jenkins, who endorses his point) goes too far when he asserts that 'we knew already' that 'Non-scepticism . . . is the practical alternative'. It needs an argument that, in deference to the sceptical paradoxes, one could not live a manageable life by the lights of some kind of instrumentalist, or Van Fraassen-style, constructive empiricist conception of the 'hypothesis' of the external world, or of a substantial past, or the existence of other minds—the solipsist might be lonely, but must he be *practically* disabled? But the objection that Pritchard is making does not depend on his over-egging it in that way. Its central point is that entitlement, in the best case, only amounts to a kind of pragmatic warrant, whereas the challenge of scepticism is to provide *epistemic* warrant.

For my present purposes, I'll take the contrast intended by Pritchard and Jenkins between properly epistemic warrant and the merely pragmatic to be illustrated by the argument of Pascal's wager, schematized in decision-theoretic tabular form as follows:

	<i>GOD EXISTS</i>	<i>THERE IS NO GOD</i>
<i>BELIEVE</i>	ETERNAL BLISS	OBLIVION
<i>DON'T BELIEVE</i>	BAD PLACE	OBLIVION

Suppose we grant that the schematized reasoning somehow shows that it is rational to believe in God.³⁸ Still, it's very intuitive that the justification thereby provided is at best

³⁷ Jenkins 2007: 27.

³⁸ In fact, of course, the entries are incomplete: what if there is indeed a God, but one whose patience is wearing thin with people whose acceptance of Him is motivated solely by Pascal-style pragmatics?!

a *prudential* one. Whereas—Pritchard and Jenkins are saying—the sceptical challenge was to show that our ordinary beliefs, based on perception, induction, and so on, are *epistemically* warranted.

That was indeed the challenge, but what exactly is its content? The terms in which it is framed are hardly pellucid. If ‘epistemically warranted’ is merely *defined* to mean: supported by evidence, then sure, neither Pascal nor Reichenbach provides an epistemic warrant. And if the sceptical challenge is rightly formulated as being to provide epistemic warrant so defined—non-question-begging evidence—for the truth of cornerstones, then any response along broadly Reichenbachian lines—and indeed anything except the adduction of such evidence—will be an evasion. But is the essential sceptical challenge rightly so formulated? The underlying question concerns what it is that is *properly intellectually disturbing* about scepticism: why we should care about the evidential isolation of the cornerstones that the sceptical paradoxes seem to highlight? It *sounds* diagnostic to say, ‘Well, we should care because the conclusion that threatens is that we lack all epistemic justification for these fundamental beliefs’. But it is a superficial diagnosis unless ‘epistemic warrant’ is not merely restricted by stipulation to what is afforded by evidence and cognitive achievement but is first associated with some independently explained value and then argued to be appropriately restricted to what is so afforded, because only evidence and cognitive achievement can provide us with that value. It is, in other words, a substantial question, not to be finessed by a favoured but unargued definition of ‘epistemic’, whether what is properly intellectually disturbing about scepticism can be adequately addressed by consideration of entitlements and the extent to which we have them.

Let’s approach the issues here by asking what it is that Pascal-type argument can actually, in the best case, deliver. What, if any, is the real difference between the kind of reasons bestowed by the possession of suitable evidence and the kind afforded by best-case Pascal? Is there really a significant distinction, sufficient to ground enough of a difference in their respective value to justify the Pritchard/Jenkins complaint?

The question can seem quite hard. One distinction I have heard proposed in conversation is that the kind of warrant generated by Pascal-type reasoning is not properly described as *reason to believe* at all—that it is, rather, *reason to want* to believe. This is not terribly convincing. Doubtless, Pascal-type reasoning, if good, does generate reason to want to believe. But can that be kept far enough apart from reason to believe to sustain the intended distinction? The thought seems plausible that whatever one has some reason to want to do, one has the very same reason to do—that an answer to the question, ‘Why should I want to do that?’ will be an equally good answer to the question: ‘Why should I do that (if I can)?’ If so, then, provided we allow that Pascal’s reasoning does indeed give one reason to want to believe in God, it gives one the same reason to believe in God.

It may be rejoined that, even if that is granted, there is still the difference that Pascal-reasoning supplies reason to believe only *by* supplying reason to want to believe, whereas proper evidence supplies reason to believe *irrespective of one’s wants*.

But this too seems not quite right. Presumably one should want to be rational and hence should want to have those beliefs that are mandated by one's evidence. And in that case good evidential reason to believe *P* will also provide reason, as far as it goes, to want to believe *P* (even if one may have other reasons for wishing not to believe it—as when the truth hurts). The simple proposal, that we can sharply separate reason to want to believe from reason to believe, and that Pascal at best supplies the former without the latter, and evidence the latter without the former, doesn't seem to hit off the intuitive distinction cleanly.

Still, what I think is the right suggestion is nearby. Even if evidence does, in conjunction with an associated background desire to respond rationally to one's evidence, provide reason to want to believe the propositions for which it is evidence, the reason it provides for believing is not—or need not be—*mediated by* that want: if someone believes *P* because they have good evidence for it, we don't need to take it that they must have reacted to the evidence in the light of a standing desire to believe what the evidence supports and by then acting on that desire. Relatedly, the formation of a belief on the rational basis of another belief does not need to be mediated by desire at all. In short, the mental acts involved in rational belief formation and management are not essentially subject to a practical syllogism—a syllogism configuring both belief and desire components. The rational explanation of action in general does need to advert to psychological factors of both kinds. But the rational explanation of the mental act of coming to a belief need refer only to the subject's experiential states and/or other beliefs. By contrast, Pascal's considerations, even were they fully cogent, provide a rationalization for belief that entirely depends upon the desires and goals of the thinker; addressed to a rational subject who is unconcerned at the prospect of eternal bliss, they are powerless to move.

This, I suggest, is the (best account of the) sense in which Pascal-type considerations provide reasons that are properly described as practical, or pragmatic. And the same is true of any broadly Reichenbachian justification of certain kinds of trusting or associated belief-forming procedures. The case will be that it is, broadly, advantageous, in the light of our goals and values, to proceed in a certain way. Evidence, by contrast, provides reasons for belief whose motivating power in no way depends upon our sense of the prospective advantages, relative to our goals and values, of that belief.

Very well. The question now, then, is why this should matter: why it should be cause for intellectual regret if purely evidential justification for certain of our trustings and belief-forming procedures is not forthcoming. I think there is a tendency to be hampered in one's thinking about this by the sneer associated in this context with the term 'pragmatic'—one always hears it as: *merely* pragmatic, as associated with resonances of opportunism and political compromise. It is good to shift the example so that goes away. So suppose that instead of eternal bliss for myself, what is in prospect, if I successfully undertake Pascal's wager, is a very considerable alleviation of Third World suffering, poverty, and disease. In that case, I think we would be happy to say that there is a strong *moral* reason—admittedly an unusual one, but no less strong for that—for undertaking the wager

and trying to bring it about that I have the appropriate belief. It is still a pragmatic reason. But the shift in the example brings out that pragmatic reasons are not a special genre of reason, to be contrasted with, for example, epistemic, moral, and prudential reasons. The key point about pragmatic reasons is that they are contingent on the goals of the agent—the rational explanation of an action performed in the light of a pragmatic reason will take the form of an ordinary practical syllogism involving belief and desire. There is therefore no good cause to deny certain kinds of pragmatic reason the title ‘epistemic.’ This will be the case where, in the slot in the structure of the reasons for an action that is to be filled by the desires of the agent, the relevant desires are focused on *epistemic* goods and goals. In this sense, to show that certain trustings and unevidenced acceptances further the attainment of epistemic goods—of truth, understanding, and the anticipation of future experience, for example—is to provide *epistemic* reason for those acceptances. There is no point in denying the title.

In sum: ‘pragmatic,’ or ‘practical’ reasons, in the sense appealed to in the Pritchard/Jenkins objection, are not really a *kind* of reasons, apt to be contrasted with epistemic reasons, in the fashion the objection, at least as originally formulated, requires. Pragmatic reasons are instrumental reasons that take their particular *genre* from the kind of goals or values, or desires, to which they are in service. An entitlement, grounded in a broadly Reichenbachian way, marks the presence of a pragmatic reason. But because the values to which it is in service are epistemic values—the maximizing of true and useful belief—it is also an epistemic reason. The question for the objector is accordingly, why isn’t that enough? There is indeed the difference I have suggested between the architecture of the rational influence of the two kinds of epistemic reason.³⁹ What we have still to see, though, is why, if it is, that difference is important: why it should be cause for dismay if, with hinges and their ilk, we have only the one kind of epistemic reason for acceptance, and not the other. This is the key issue. I’ll return to it in Section 11.7.

³⁹ To avoid misunderstanding, let me emphasize that I here take no stand on the question central to Thomas Kelly’s 2003: the question whether epistemic reasons *in general* should be conceived as pragmatic—that is, as means-end, or instrumental reasons. Kelly argues that it is a mistake to construe epistemic rationality as a special case of instrumental rationality. Jenkins, for her part, is quick to entwine her critique of my proposal of the notion of Entitlement of Cognitive Project with the issues concerning ‘epistemic consequentialism’—characterized as the idea that ‘the epistemic status of an attitude is determined by the epistemic value of its consequences *compared to* the epistemic value of the consequences of the alternatives’ (Jenkins 2007: n. 10). But the contention that, when harnessed to expected epistemic utility, a warrant grounded in Reichenbachian considerations may properly be termed epistemic, entails nothing about the question whether *all* epistemic reason is similarly instrumental or pragmatic in character. More specifically, there are two different questions here, it seems to me: whether the *rational force* of reasons for belief bestowed by evidence is grounded in the instrumental value of evidence towards the attainment of other goals, like truth and knowledge; and whether the *motivating power* of evidence works through and is contingent on the goals of the subject. I have taken a negative view on the second in distinguishing evidential from pragmatic reasons in the manner suggested in the text. But I have taken no view on the first—the issue I take Kelly and Jenkins to have in mind—and I do not need to (though obviously it would further empower the contention that Reichenbachian warrants, grounded in the enhancement of expected epistemic utility, are properly regarded as epistemic warrants if evidential warrant *itself* is so grounded!).

11.6. Is Entitlement A Priori?

It seems reasonable to suppose that if an appeal to entitlement is to provide means to address the essentially a priori challenge of sceptical paradox, then the considerations that ground entitlements must be a priori too. Michael Williams⁴⁰ objects that this constraint cannot be met. It cannot be met precisely because, at least on the Reichenbachian model, the non-evidential warrants that are grounded in entitlements are indeed pragmatic warrants, whose validity is accordingly contingent on our epistemic aims and values. Since it is an a posteriori question what our epistemic aims and values are, the warrants, if any, issuing from the considerations that are supposed to ground entitlements cannot be resolutely a priori. Where Hume postulated contingent if unalterable doxastic propensities and limitations, the defender of entitlement—Williams thinks—must postulate contingencies of human aims and values, knowable only a posteriori. The contingencies concerned may be unquestioned. But the fact will be that empirical information will be being smuggled into the grounding of entitlements; and there will then be a question, in a context where other-minds scepticism, for example, is on the table, why that is not a foul.

I think it is not a foul for the following reason. Even if entitlement is construed as a pragmatic (though still epistemic) ground for trust in certain propositions—a ground that depends upon contingencies about what we value or aim for in enquiry that are knowable only a posteriori—it does not follow that the considerations which disclose entitlements will not be available at a purely reflective level of consideration: the ‘level of philosophy’, as Williams likes to say. For the sting of sceptical argument itself is conditional upon those same values. If scepticism poses a challenge, it is a challenge to make out how it can be rational to accept large swathes of our common-sense beliefs on the grounds on which we actually do accept them. The challenge is posed in a context in which it is taken for granted that we value truth, and value evidence of truth, and are consequently properly disturbed by argument that seemingly calls into question the propriety of our evidence, the coherence of the standards to which we subject it, or its very existence. Epistemological discussion at the ‘level of philosophy’ operates under the aegis of these assumptions. If, *per impossibile*, someone was to produce, by way of a response to scepticism, a genuinely forceful evidential argument for believing that there is an external material world, it would be bizarre to respond that the force of the argument was qualified by its dependence on the contingencies of our epistemic values that, in context of addressing scepticism, could not legitimately be presupposed. Philosophy is already up to its neck in presuppositions about the values that we allow to govern enquiry. If that means that philosophical discussion about paradoxes of enquiry is not purely a priori, then it was not true to begin with that the grounding of entitlements, if to be accomplished in a way relevant to addressing scepticism, has to be demonstrated purely a priori. Rather, in attempting to respond to scepticism, we

⁴⁰ Williams 2012.

may legitimately appropriate all the resources involved both in the setting of the paradoxes, and in the concerns they generate.

11.7. Real Entitlements

Let it be that Reichenbachian considerations can in principle generate something properly described as a non-evidential but still epistemic warrant to accept a certain proposition—a warrant connected to expected epistemic utility. Our question was: why, if that is all we have, are we worse off than if we had an evidential justification for the same proposition? After all, it might be contended, reliance upon evidence is itself rational, arguably, only in the light of the expected epistemic utility of a policy doing so—in the light of the expectation that we thereby enhance our chances of believing the truth and avoiding error.⁴¹

Let us construct a case as close as possible to Pascal. Suppose that a mischievous but broadly benevolent Demon offers to teach me a range of undreamt-of and very powerful truth-productive methods of mathematical enquiry if I can but bring myself to believe Goldbach's Conjecture while we still lack any proof of it. Even if we allow that the Demon's offer gives me powerful epistemic reason to believe the Conjecture, the feeling is apt to remain that there will still be something amiss with the belief. The same would be true if the benefits of believing were moral. I'd be morally justified, we might say, in forming the belief if I possibly could, but there would still be something criticizable, or off-colour, about doing so.⁴² The essence of the complaint that Pritchard and Jenkins are making, it seems to me, is that the kind of warrant for an unevidenced proposition generated by a Reichenbachian strategy, even if properly described as epistemic, will still be warrant that involves *compromise*—compromise of our (other) standing epistemic values: the very values tugged at when we confront the sceptical paradoxes. By contrast, the kind of warrant that we should like a satisfactory theory of epistemic entitlement to provide us with should involve no such epistemic compromise. The reason I am given by the Demon's offer to believe Goldbach's Conjecture is one that that *competes with*—and perhaps overpowers—my countervailing reasons *not* to believe in Goldbach's Conjecture: viz. the fact that I lack any proof of it, or any other kind of corroborative evidence. Epistemic reason of the kind delivered by Pascalian or Reichenbachian reasoning, when it amounts to sufficient reason, does so, we may well feel, by overcoming the basic, orthodox reason, supplied by our lack of evidence for the targeted proposition, for not believing it. In this respect it stands comparison with overwhelming moral reason for doing something that is independently morally regrettable—for punishing the innocent in certain circumstances, perhaps, or sacrificing the proverbial Fat Man on the railway track.

⁴¹ But see n. 39.

⁴² Cf. Jenkins 2007: 37.

Another way to put the objection is to say that, to the extent that the notion has so far been explained, no reason has so far been given to deny that an *entitlement* to believe Goldbach's Conjecture is generated when the Demon makes his bizarre offer. I agree. A better theory of entitlement than anything provided by Reichenbachian argument ought to give us such a reason. And a good theory of entitlement, while foreseeably allowing Reichenbachian reasonings to underwrite warrant for certain acceptances, should therefore augment them by other considerations, so far un-indicated, why responding positively to such considerations, or other forms of non-evidential warrant, in circumstances where one has no evidential warrant for certain targeted propositions, need actually involve no compromise of others of our epistemic values, when clear-headedly conceived. The propositions that pass *that* test will be the *real* entitlements.

In sum: a real epistemic entitlement must meet two conditions. First, there must be good reason why a trusting, unvidenced acceptance of the proposition in question may be expected to enhance expected epistemic utility. That is what Reichenbachian considerations, for example, where applicable, may contrive to provide. But that alone fails to distinguish the case from the offer of the Goldbach Demon. The second requirement is that a trusting acceptance of the proposition in question should involve no *compromise* of legitimate epistemic standards—that it should not merely be the better of two alternatives each of which comes short, at least to some degree, of fulfilling legitimate epistemic values. And of course the simple, obstructive thought here is that the acceptance of any proposition as true without evidence or relevant epistemic achievement *always* involves compromise, since it is irrational to take any proposition purely on trust. Isn't that merely the proper, scientific outlook?

11.8. Against the Evidential Ideal

Scientific perhaps. But proper? No. Rather, the obstructive thought is exactly what needs to be dislodged. Of course it is often, perhaps normally, irrational to confidently assume or take the truth of propositions on trust. But the basic insight behind the entitlement project—Wittgenstein's insight, I believe—is that *all* reflective enquiry, and all reflective cognitive accomplishment, is essentially situated in trusting acceptances, some general, others specific to the particular context of enquiry, for which we lack evidence. This is not a shortcoming, a lapse which, though unavoidable, is nevertheless regrettable. It is in the *nature* of rational reflective inquiry that this should be so. There is no coherent model of the reflective acquisition of knowledge and warranted belief in which it is not so. There is simply no such thing as the epistemically clean slate; and the notion of a kind of dawn of enquiry from which ideal epistemic advance would involve only the addition of beliefs fully attested by verification or good but defeasible evidence is incoherent. The accumulation of evidential reason to believe is possible only within the framework set by trusting acceptances. Without them, we forfeit, in particular, our detailed conceptions of evidential relevance—of why, for example,

perceptual experience bears on how matters are configured around us in the external world. The entitlement project is to characterize the various roles of the acceptances, the specific aspects of our enquiries that they sustain and rationalize. It remains to be seen in what degree of convincing detail the project can be executed. But the basic point must be that there is no way things could be otherwise: without trusting, there is no reflective enquiry.

I say, 'no reflective enquiry'. There is a point here that marks a departure from Wittgenstein—at least if we may go by the letter of the notes in *On Certainty*. Wittgenstein repeatedly writes as though trust, or uncritical acceptance, were integral to one's epistemic life. My suggestion is that they are integral to epistemic life within the sphere of what Sosa⁴³ calls 'reflective knowledge'—or what I have been calling 'the level of claims'. The quotidian acquisition of earned perceptual beliefs, for example, or of 'animal knowledge' in Sosa's sense, need involve no placing of trust in the effective functioning of one's perceptual capacities—a small child need not even have any concept of perceptual faculties as such, yet can still, of course, be forming perceptually justified beliefs all the time. Authenticity-conditions, and hinges in general, only surface into our thinking as commitments at the level of claims to knowledge and warranted belief.

Sceptical paradox works by canvassing what seems like an intuitively rational ideal—that an epistemically rational thinker will apportion her beliefs to the evidence and so, in particular, will accept nothing for which she has *no* evidence or, at the level of reflection, for which she has nothing that she has good reason to take to be evidence. Call this the Evidential Ideal. The paradoxical arguments then work by developing apparent challenges to our right to regard large species of normally accepted evidence as in good standing or genuinely probative. But the theorist of entitlement is contending that the Evidential Ideal is incoherent. There is no such thing as apportioning one's beliefs to the evidence unless some things are accepted without evidence, since it is only in the context of other acceptances that we can have determinate conceptions of what is evidence for what, or of what our cognitive powers are. To lament the fact that reflective enquiry rests upon a repository of trust is thus a bit like lamenting the fact that not everyone can be at the front of a queue, or that not everyone can be in fiscal credit. Being at the front of the queue brings determinate advantages for those there; being in credit gives powers and freedoms to those in credit. There is a superficially coherent thought that it would be nice if the advantages, powers, and freedoms concerned were universally shared. But it is only superficially coherent. There is a similarly superficially coherent thought that it would be nice if all our beliefs rested upon secure evidential foundations. But again, it is only superficially coherent. And there is certainly no coherent conception of epistemic rationality which views unevidenced trusting as *per se* a lapse, or in our terminology above, a compromise of it. Epistemic rationality better amount to a coherent set of standards governing all enquiry. There is no such coherent set of standards that has that consequence.

⁴³ Sosa 2007 and 2009.

I have encountered in discussion⁴⁴ the response that, even if all that is true, it doesn't really change matters. Let it be that reflective enquiry necessarily rests upon a framework of specific and general trustings that are not themselves the products of enquiry. Still, a *rationale* is thereby provided for the relevant trustings only if enquiry is valued. We have the option of dropping out, or dropping dead. But I think that to accept that enquiry is not forced upon us is quite consistent with the entitlement proposal's being, at least potentially, a fully satisfying response to scepticism. If I am charged to give you a moral justification for a certain course of action, and proceed to offer considerations which demonstrate, for example, the injustice of the only available alternative, it would be strange to reply, 'Well, that only works if you presume that I am concerned about justice'. In the same way, I think we have given a properly epistemic justification of certain aspects of our practices if it is demonstrated that they enhance the attainment of epistemic value, and involve no compromise of epistemic rationality. And the latter is secured if, as I have just been stressing, these aspects of practice are integral to enquiry itself. A demonstration that that is so will be philosophical; it will respond to scepticism within the rules of debate set by scepticism itself. And it will be an epistemic justification. It is another question whether we actually have, or must have, the epistemic values to which enquiry is harnessed, or whether we value enquiry itself. It is not an issue which *epistemic* justification has to tackle. Nor does a moral justification have somehow to argue that we have no option but to value the moral goods to which it appeals.

One last wriggle. Some sceptical person might say that the real import of the consideration that all fully reflective enquiry must involve unevidenced acceptances is not that the Evidential Ideal is incoherent but that it can be satisfied only by abstaining from enquiry altogether. Actually, I am inclined to grant this. But just for that reason, we should wonder why notions of rational enquiry and rational acceptance with so suicidal a consequence should be held to be any kind of *ideal*. A morality that had the consequence that one should refrain from all action lest one inadvertently do harm would fall down on the basic requirement that morality serve to promote ethically valuable—just, beneficial, etc.—action. If moral principles seemed intuitively good to us that had so negative an upshot, we would do well to learn the lesson that such principles need to be rethought and refashioned into something more to the purpose. It is no different in the epistemic sphere. An account is owed in the first place of why an exceptionless prohibition on unevidenced acceptance should seem appropriate. But even if a powerful such account could be provided, the result of accepting such a principle is merely going to be an unworkable ethics of belief management. I do not think that the Evidential Ideal is actually a considered part of anything worth regarding as our intuitive notion of epistemic rationality. But if someone disagrees, then they should conclude that they

⁴⁴ And in print: I take the train of thought following to be that at pp. 28–31 of Jenkins 2007.

would do well to revise their notion of epistemic rationality. Entitlements, in the general sense canvassed here, will be a part of any workable revised system.

11.9. Concluding: The Problem of Demarcation

Strategic Entitlement—entitlement grounded in broadly Reichenbachian, or decision-theoretic considerations—was one of four possible models of entitlement outlined in Wright (2004a). In this paper I have been primarily concerned to revise the negative assessment of this model that I sustained in the previous paper, and to explain how a proponent of it, or any of the other models, may respond to the various objections we have considered. I have not here further pursued those other models of entitlement—entitlements of Cognitive Project, of Rational Deliberation, and of Substance. And I have not further encroached at all on the major question of the *demarcation* of entitlement: even if some notion of warranted but unevidenced acceptance does have an ineliminable part to play in any feasible methodology of enquiry and belief management, the question remains, how big a part, and what are its limits? The point has not gone away that it is not in general, or even usually, consistent with responsible belief management to accept things without evidence or relevant cognitive achievement. What are the principles that determine when one may do so and when one ought not? How do we distinguish the genuine entitlements from the prejudices, mere assumptions, and *idées fixes*? No less important than trying to delimit by what principles we may be rationally entitled to certain trustings is the project of determining when we are *not*, that is, when absence of evidence does indeed defeat rational acceptance.

This is, of course, an absolutely crucial issue. It presents, in my judgement, perhaps the most major challenge remaining to the theorist of entitlement. But it is something that here I will have to reserve for another occasion.

References

- Altschul, Jon (2011). 'Epistemic Entitlement.' *Internet Encyclopaedia of Philosophy* at <<http://www.iep.utm.edu/ep-en/>>.
- Burge, T. (1993). 'Content Preservation.' *Philosophical Review*, 102 (4): 457–88.
- (1996). 'Our Entitlement to Self-knowledge.' *Proceedings of the Aristotelian Society*, 96 (1): 91–116.
- (2003). 'Perceptual Entitlement.' *Philosophy and Phenomenological Research*, 67 (3): 503–48.
- Coliva, A. (ed.) (2012). *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*. Oxford: Oxford University Press.
- Davies, M. (2004). 'Epistemic Entitlement, Warrant Transmission and Easy Knowledge.' *Aristotelian Society*, Supplementary Vol. 78 (1): 213–45.
- Dretske, F. (2000). 'Entitlement: Epistemic Rights without Epistemic Duties?' *Philosophy and Phenomenological Research*, 60 (3): 591–606.

- Dretske, F. (2005). 'The Case against Closure.' In M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology*, 13–26. Oxford: Blackwell.
- Hawthorne, J. (2004) *Knowledge and Lotteries*. Oxford: Clarendon Press
- (2005). 'The Case for Closure,' In M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology*, 26–43. Oxford: Blackwell.
- Holliday, W. (forthcoming) 'Fallibilism and Multiple Paths to Knowledge.' In T. Szabo Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, volume 5. Oxford: Oxford University Press
- Jenkins, C. (2007). 'Entitlement and Rationality.' *Synthese*, 157 (1): 25–45.
- Kelly, T. (2003). 'Epistemic Rationality as Instrumental Rationality: A Critique.' *Philosophy and Phenomenological Research*, 66 (3): 612–40.
- Moyal-Sharrock, D. and W. H. Brenner (eds.) (2007). *Readings of Wittgenstein's On Certainty*. Basingstoke: Palgrave Macmillan.
- Peacocke, C. (2004). *The Realm of Reason*. Oxford: Oxford University Press.
- Pritchard, D. (2007). 'Wittgenstein's *On Certainty* and Contemporary Anti-Scepticism.' In D. Moyal-Sharrock and W. H. Brenner (eds.), *Readings of Wittgenstein's On Certainty*, 189–224. Basingstoke: Palgrave Macmillan.
- Pryor, J. (2000). 'The Skeptic and the Dogmatist.' *Noûs*, 34: 517–49.
- (2013). 'Problems for Credulism.' In C. Tucker (ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, 89–131. New York: Oxford University Press.
- Reichenbach, H. (1938). *Experience and Prediction*. Chicago: University of Chicago Press.
- (1949). *The Theory of Probability*, 2nd edn., tr. E. H. Hutten and M. Reichenbach. Berkeley: University of California Press.
- (1968). *The Rise of Scientific Philosophy*. Berkeley and Los Angeles: University of California Press.
- Russell, B. (1921). *The Analysis of Mind*. London: G. Allen & Unwin.
- Salmon, W. (1991). 'Hans Reichenbach's Vindication of Induction.' *Erkenntnis*, 35: 99–122.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Vol. I. New York: Oxford University Press.
- (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Vol. II. New York: Oxford University Press.
- Steup, M. and E. Sosa (eds.) (2005). *Contemporary Debates in Epistemology*. Oxford: Blackwell.
- Tucker, C. (ed.) (2013). *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. New York: Oxford University Press.
- Williams, M. (2012). 'Wright against the Sceptics.' in A. Coliva (ed.), *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*, 352–75. Oxford: Oxford University Press.
- Wittgenstein, L. (1969) *On Certainty*, ed. G. E. M. Anscombe and G. H. von Wright, tr. D. Paul and G. E. M. Anscombe. Oxford: Blackwell.
- Wright, C. (2004a). 'On Epistemic Entitlement: Warrant for Nothing (and Foundations for Free?)' *Aristotelian Society*, Supplementary Vol. 78 (1): 167–212.
- (2004b). 'Intuition, Entitlement and the Epistemology of Logical Laws.' *Dialectica*, 58: 155–75.
- (2007). 'The Perils of Dogmatism.' In S. Nuccetelli and G. Seay (eds.), *Themes from G. E. Moore: New Essays in Epistemology and Ethics*, 25–48. Oxford: Oxford University Press.

- (2008). 'Internal–External: Doxastic Norms and the Defusing of Skeptical Paradox.' *Journal of Philosophy*, 105 (9): 501–17.
- (2012). 'Replies Part IV: Warrant Transmission and Entitlement.' In A. Coliva (ed.), *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*, 451–83. Oxford: Oxford University Press.
- Yablo, S. (forthcoming) *Aboutness*. Princeton: Princeton University Press.