

Entitlement and rationality

C.S. Jenkins

Received: 2 November 2005 / Accepted: 22 March 2006 / Published online: 31 January 2007
© Springer Science+Business Media B.V. 2007

Abstract This paper takes the form of a critical discussion of Crispin Wright's notion of entitlement of cognitive project. I examine various strategies for defending the claim that entitlement can make acceptance of a proposition epistemically rational, including one which appeals to epistemic consequentialism. Ultimately, I argue, none of these strategies is successful, but the attempt to isolate points of disagreement with Wright issues in some positive proposals as to how an epistemic consequentialist should characterize epistemic rationality.

Keywords Entitlement; Epistemology; Epistemic Rationality

Introduction

Crispin Wright has recently suggested (Wright, 2004) that one can possess a form of warrant for a proposition p which is 'beyond rational reproach even though [one] can point to no cognitive accomplishment . . . whose upshot could reasonably be contended to be that [one] had come to know p , or had succeeded in getting evidence justifying p ', where 'evidence' is to be understood 'in the broadest sense, encompassing both a priori and empirical considerations' (pp. 174–175).

In particular, Wright argues that one has 'entitlement of cognitive project' in connection with certain propositions. To qualify, a proposition p must be a 'presupposition' of a cognitive project (i.e. a proposition doubt about which would rationally commit us to doubting the significance or competence of the project), and that project must be one that is so important for us to pursue that its 'failure would at least be no worse than the costs of not executing it, and its success would be better'. Moreover, two further conditions must be met:

- (i) We have no sufficient evidence to believe that p is untrue, and

C.S. Jenkins (✉)
Arche AHRC Research Centre, Philosophical and Anthropological Studies,
University of St. Andrews, 17–19, College Street, St. Andrews, KY16 9AL Fife, UK
e-mail: caroline.jenkins@st-andrews.ac.uk

- (ii) The attempt to justify *p* would involve further presuppositions in turn of no more secure a prior standing . . . and so on without limit; so that someone pursuing the relevant enquiry who accepted that there is nevertheless an onus to justify *p* would implicitly undertake a commitment to an infinite regress of justificatory projects, each concerned to vindicate the presuppositions of its predecessor (pp. 191–192).

To illustrate, it seems that according to Wright I have an entitlement to trust that:

S *My sensory apparatus is (for the most part) putting me in touch with the external world.*¹

S is a presupposition of many crucial cognitive projects, and I am prepared to grant (at least for the sake of argument) that these projects are so important that their failure would be no worse than their non-execution (and of course, I allow that their success would be better).² Moreover, I have no sufficient evidence to believe that *S* is untrue, and any attempt to justify *S* would involve further presuppositions of the kind described in condition (ii).

The significance Wright attaches to the notion of entitlement of cognitive project is considerable: he would take the claim that we are entitled to trust that *S* to provide the basis for a response a certain form of sceptical argument. Entitlement to trust in some such proposition, he says, would ‘immediately [empower] us to dismiss the various scenarios of cognitive dislocation—dreams, sustained hallucination, envatment and so on—which are the stock-in-trade of Cartesian scepticism’, and if this can be done then ‘the Cartesian sceptical argument . . . is nipped in the bud’ (p. 195).

Wright admits that what he has in mind is a ‘sceptical solution’ to the worries raised by Cartesian scepticism. We do not recover any claim to *know S*. His aim is to show that when we trust that *S* is true, when we rely on *S*, ‘there is no irrationality, or capriciousness, in our proceeding in the ways we do—that we are warranted in so proceeding’, albeit warranted in a non-evidential way (p. 206).

Before we can assess this claim, we need to get clear as to what is meant. Wright uses the word ‘rational’ and its cognates frequently in his discussion; entitlement is supposed to be something which turns mere acceptance into ‘rational trust’. But he never mentions (even to subvert it) the traditional distinction between *epistemic* and *practical* rationality. The distinction is easy to motivate; it can, for instance, be rational in the practical sense to adopt some propositional attitude just because you know that if you do so you’ll be given chocolate (which you like and desire), and there will be no consequences that you do not desire. But it clearly wouldn’t be *epistemically* rational to adopt that attitude just because you knew it would have this result. For now, however, to avoid begging any important questions against Wright, I shan’t say any more about what epistemic rationality is, but shall leave the distinction between practical and epistemic rationality at the level of intuitions.

Whatever the distinction amounts to, however, it is important in the current context. Wright’s silence on the point leaves him open to interpretation along the lines of

¹ NB Wright does not think that entitlement of cognitive project can help us with ontological questions concerning the existence and nature of the external world, but that if ‘a certain conception of our cognitive powers and their sphere of operation’ is granted, we may be entitled to trust that ‘those powers are functioning effectively in conducive circumstances’ (p. 197).

² The sense in which failure would be ‘no worse’ and success would be ‘better’ will need careful discussion below.

Pritchard (2005, §3), who takes the upshot of Wright's discussion to be that '[n]on-scepticism is defended on the grounds that it is the practical alternative'. And as Pritchard says, 'we knew *that* already.' 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.

Moreover, Wright's use of the term 'warranted' to describe the status he hopes to establish for propositions like S strongly suggests that what is at issue is the epistemic status of these propositions—the epistemic rationality of the 'rational trust' we place in them. Hence both charity and textual evidence support the hypothesis that Wright is trying to establish that entitlement is what turns mere acceptance into *epistemically* rational trust.³

Doing without an appeal to dominance

At first blush, the consideration to which Wright seems to be appealing in order to argue that entitlement of cognitive project is a form of epistemic warrant for trusting that S is:

W *We ought to trust that S, if we want to do well epistemically, since trusting that S cannot make our epistemic situation any worse and may make it much better.*

Notice that, although the definition on p. 191 implies that we have entitlement of cognitive project to *any* presupposition of *any* project, it is clearly intended (see p. 192) that a genuine entitlement is created only when we are dealing with projects which are such that their failure would at least be no worse than the costs of non-execution and their success would be better. (Note also that, assuming Wright's goal is not merely a defence of the practical rationality of trusting that S, I have charitably interpreted

³ Wright discusses four species of entitlement in his paper. My reason for focussing on entitlement of cognitive project is that, in my opinion, it is the one which holds the most serious promise of establishing the epistemic rationality of the propositions to which we are entitled.

Wright as meaning ‘epistemically better’ and ‘epistemically worse’ by ‘better’ and ‘worse’ in this passage. More on this point in the next section.)

The felt need thus to restrict the conditions under which we have entitlement of cognitive project strongly suggests that an appeal is being made to something like W. However, any appeal to W might be thought to make the notion of entitlement of cognitive project sound rather like the notion of ‘strategic’ entitlement which Wright discusses in an earlier section of his paper. Strategic entitlement explicitly appeals to a form of game-theoretic dominant-strategy reasoning (of roughly the kind exemplified by Reichenbach (1938) on the problem of induction). To say that something is a dominant strategy is to say that in every situation it does at least as well the alternatives, and in some situations it does better. We secure strategic entitlement to a proposition just in case acting as if one had a justified belief in that proposition cannot hinder one’s purposes and may promote them (p. 183). Entitlement of cognitive project, however, is supposed to be ‘rather different’ from strategic entitlement.⁴

Moreover, in conversation Wright has suggested that he did not intend to appeal to W. (Although I will examine the prospects for doing so in the next section, I do not imagine that Wright himself would find that option attractive.) So is there any other way to read the claim made for entitlement of cognitive project, which avoids the appeal to dominance reasoning? One way might be to read Wright as making a point akin to the ethical claim that one cannot be morally blamed for doing something one is forced to do. Certain turns of phrase support this sort of interpretation; for instance, Wright begins his section on entitlement of cognitive project by discussing two passages from Wittgenstein’s *On Certainty*, taking the thrust of them to be that ‘one *cannot but* take certain . . . things for granted’. And it is the fact that checking all the presuppositions of one’s cognitive projects is *impossible* that motivates Wright’s claim that a ‘proper concept’ of warrant should not make such demands (p. 190).

If I cannot but kill someone, I should not be morally blamed when I do kill him. Similarly, perhaps, if it is impossible to avoid trusting in some proposition, then I should not be epistemically blamed (accused of epistemic irrationality) when I do so trust. But any analogy Wright intends to draw can’t be quite that simple. For a start, all he is claiming is that it is impossible to avoid trusting in some propositions *if* we

⁴ It is not clear (to me) exactly what all the differences are supposed to be. For one thing, early on it looks as if strategic entitlement to a proposition only establishes the rationality of *acting* (in some—perhaps all—situations and contexts) *as if* we had a justified belief that that proposition were true, without giving any grounds for ‘subjective confidence’ in its truth (p. 185; see also p. 192). Entitlement of cognitive project, by contrast, is meant to be entitlement to trust in a way that excludes agnosticism about the proposition’s truth. If this were the correct reading, any ‘rational authority’ generated by strategic entitlement would seem to be practical, rather than epistemic, since it would be applicable to actions in general, and not only to cognitive acts. (This might supply an explanation of why, according to Wright, the notion of strategic entitlement is only helpful where *methodological* scepticism is at issue.)

Later, however, it is argued that strategic entitlement must be entitlement to trust (in the agnosticism-unfriendly way), if it is to ‘underwrite a policy of *belief-formation* (rather than merely non-doxastic forms of action)’ (pp. 193–194, emphasis in the original). But no argument is given that strategic entitlement *could* be (or give rise to) entitlement to trust in this sense. So it is not clear whether Wright thinks it could, or whether this is, in effect, the basis for a modus tollens argument against the possibility of using strategic entitlement to underwrite policies of belief-formation. One suspects the former, but we await an argument.

want to engage seriously in the cognitive projects for which they are presuppositions.⁵ ('Seriously' here is my shorthand for 'while rationally⁶ refraining from any doubt about their significance or competence'). The obligation to trust, then, is conditional upon our wanting to engage seriously in those projects. Thus we could respond to the argument under consideration here by saying that the epistemically rational (non-blameworthy) thing to do would simply be to refrain from engaging seriously in the projects in question. By analogy, I cannot be morally excused from killing someone just because I had to kill him if I was to get the job I wanted. The non-blameworthy thing to do in that situation would be to refrain from getting the job.

You might think I could be morally excused from killing someone just because I had to kill him if I was to get the job, provided that I was for some reason *morally obliged* to get that job. Is there a corresponding possibility that there are some cognitive projects that we are *epistemically* obliged to engage in seriously? That is, are there projects such that one won't count as epistemically rational unless one pursues them seriously? If there are, this might seem to get us out of the fix. For maybe I can be epistemically excused of apparent irrationality if I accept a proposition which I have to accept in order to engage seriously in a cognitive project which it is epistemically obligatory to engage seriously in.

Leaving aside the concern that it seems difficult to give any examples of such projects,⁷ we can still take issue with the line of thought under consideration. Suppose I

⁵ The refinement collapses, of course, if we cannot but engage in those projects. But it is no part of Wright's argument—and an argument would certainly be needed—that there are any cognitive projects which we cannot but engage in (where that does not amount to saying that these projects are 'epistemically obligatory'; a discussion of the latter possibility follows shortly).

⁶ NB The word 'rationally' does important work here: its inclusion is necessary if I am to set up the dilemma I outline on p. 7 below. So it is worth stressing that the inclusion of the word 'rationally' is in line with Wright's intentions. Trust in the presuppositions of a project is not necessary in order for us to refrain from doubt about the significance and competence of the project, for one could *irrationally* refrain from doubting the significance or competence of the projects while continuing to doubt/distrust their presuppositions. What requires trust in the presupposition is the pursuit of a project while *rationally* refraining from doubt about its significance or competence. You might think it was possible to skirt this issue without including the word 'rationally' in the definition of 'seriously'. Instead, we could define 'serious' pursuit as pursuit where we refrain from doubt about the significance and competence of the project *because* we don't doubt the project's presuppositions. But then it's trivial that trust in the presuppositions is necessary for serious pursuit of the project. It we take this route, Wright's claim that trust in the presuppositions of a project is necessary if the project is to be pursued seriously amounts to the claim that trust in the presuppositions is necessary for something which is defined as including trust in the presuppositions. Correspondingly, his claim that trust in the presuppositions is warranted because such trust is necessary for serious pursuit of certain projects becomes nothing more than a claim that trust in the presuppositions is warranted because such trust is necessary for [trusting in the propositions and Φ -ing]. However important [trusting in the propositions and Φ -ing] might be, if we pursue this line we won't be able to construct an argument which has a chance of being more persuasive than an argument to the effect that trust in the presuppositions is warranted because it is necessary for trust in the presuppositions.

⁷ Couldn't there be a purely deductive rational thinker, who reasons using the principles of logic but nothing else, and a purely abductive rational thinker, who reasons using principles of abduction but nothing else? If two such thinkers are possible, then it is not clear that there is *any* project serious pursuit of which is a necessary condition for being epistemically rational.

One might doubt whether this is a genuine possibility, because one might suspect that *given the kind of beings we are and the kind of world we're in* we are epistemically obliged to engage in serious pursuit of some empirical projects. (I am indebted to an anonymous referee for this suggestion.) But I'm not sure how much weight to place on consideration of the kind of beings we are, for if we can be epistemically rational without seriously pursuing any empirical projects, even if that would turn us into another kind of being, then I don't think we are epistemically compelled seriously to pursue any

am in a very unfortunate position, whereby there is *nothing* I can rationally do that will enable me to engage seriously in the cognitive projects in question. In the envisaged situation, I could only so engage by trusting that some proposition P is true, but it is not epistemically rational to have that sort of trust in P.

According to the view under consideration here, this situation is impossible. According to that view, it *is* rational to trust that P in these circumstances, *just because* if I do not do so I cannot engage seriously in certain cognitive projects which it is epistemically essential to engage in seriously.

But that doesn't seem right; the situation I described does seem to be possible. It may be impossible to have rational trust in the significance and competence of epistemically obligatory projects, because it is impossible to have rational trust in P. There seems to be a disanalogy with the ethical case here. Saying that something is morally obligatory is often taken as implying that that thing is possible for us. But saying that something is epistemically obligatory does not seem to be like that. There is something strange about saying that Kate is morally obliged to do X, although she cannot do X. But there is nothing comparably strange about saying that Kate must do X in order to be epistemically rational, although she cannot do X.

Wright seems to disagree. He says that a 'proper concept' of warrant will not make impossible demands on us. Here we run into an exegetical subtlety: maybe Wright means to suggest that our existing concept of warrant does not make such demands on us. Or maybe he means that, *if* our existing concept of warrant makes such demands, then that concept should be *replaced* with a new one which does not. If the former, then he is simply denying the intuition which motivates the objection. While there is not much I can say in reply (except to rephrase the objection as a dilemma, as I shall do in a moment), I suspect that it is an intuition many will share.

On the other hand, if the second reading is the correct interpretation of Wright, he might be accused of changing the subject. He is no longer talking about epistemic warrant or epistemic rationality as we know them, but some new concept, warrant_w, so specified as to *guarantee* that it is possible for us to have warrant_w for the presuppositions of our cognitive projects (and hence for the many everyday propositions we believe as a result of responsibly executing those projects).

This type of strategy is not obviously misguided: perhaps a case could be made that the new concept is better suited to the purposes for which we want a concept of

Footnote 7 continued

empirical projects. And again, I cannot see (without further argument) how the kind of world we're in could place *epistemic* obligations on us to pursue empirical projects, as opposed to merely making those projects practically indispensable.

Still, even if there is no project serious pursuit of which is epistemically obligatory, surely the serious pursuit of *some* project is necessary if we are to count as rational. One suggestion which a defender of entitlement of cognitive project might develop is that the necessity of seriously pursuing *some* cognitive project if one is to count as epistemically rational generates epistemic warrant for the presuppositions of whichever basic projects one settles upon. By analogy, you might think someone can be excused of moral blame for doing X when she is morally obliged to do one of X, Y and Z but no particular one of these actions is morally obligatory for her. More needs to be said, though, since there may be some reasons to prefer Y over X which our subject can be morally blamed for ignoring. Similarly, more needs to be said in the epistemic case, since we will want to rule out certain crazy choices of basic projects. It is not clear how to go about this, though, without bringing in more theoretical resources. But if we do not try to rule out crazy choices, we seem to end up with a proposal sharing many of the unappealing qualities of Field's (2000) 'epistemological evaluativism' (which, interestingly, is motivated in a similar way to Wright's claims about entitlement, i.e. through the impossibility of giving a (non-circular) justification for basic principles—see Field 2000, p. 120ff.).

warrant than our existing concept. But it hard to see what sort of case could be made in this instance. The advantages of the new concept—its guarantee that we will be able to meet the standards set for possession of warrant for propositions like *S*—seem to be precisely the advantages of theft over honest toil.⁸

In any case, I don't think this sort of manoeuvre will help, because we can present the original objection in the form of a dilemma. On the one hand, if it is allowed that some things can be epistemically obligatory without being possible for us, the argument is in trouble as we saw above. But on the other hand, if it is maintained that anything which is epistemically obligatory is possible for us, then it will be difficult to claim, without begging the question, that any cognitive projects *T* are such that it is epistemically obligatory to pursue them seriously. For under the new assumption, if it is epistemically obligatory to pursue *T* seriously, then it is possible for us to do so. That is, it is possible for us to pursue *T* while rationally refraining from any doubt as to *T*'s significance and competence. And in order so to refrain we obviously need to be able to trust, rationally, that the presuppositions of *T* are true. But that's what we're supposed to be *showing*—we can't just assume something which obviously implies that it is true.

An argument from dominance

It seems, then, that the best chance we have of arguing that entitlement of cognitive project generates epistemically rational trust will, as initially suggested, rest upon the thought that:

W *We ought to trust that S, if we want to do well epistemically, since trusting that S cannot make our epistemic situation any worse and may make it much better.*

⁸ Besides which, some caution is needed with Wright's positive proposal that warrant (on this reading: warrant_W) is 'acquired whenever investigation is undertaken in a *fully responsible* manner', where 'full epistemic responsibility cannot, *per impossibile*, involve an investigation of every presupposition whose falsity would defeat the claim to have acquired a warrant' (p. 191, emphasis in the original). How are we to understand the idea of an investigation's being conducted 'fully responsibly'? One worry lurking here is that (at least for all the current considerations show) for any degree of responsibility (i.e. checking of presuppositions) it is always possible to increase responsibility by checking one more presupposition. This raises the question whether there is any *particular* presupposition *P* such that it is permissible (fully responsible) to accept *P* without checking. For there is no particular proposition such that *it* cannot be checked, and hence no proposition such that responsibility can't be increased by checking it. If there were some proposition like that then failure to check it would not interfere with our being 'fully responsible' in Wright's sense, since full responsibility in Wright's sense is attainable. But there isn't any proposition like that.

This point is structurally similar to that raised in fn. 7 above, and might be addressed in a similar way, i.e. by arguing that the need to stop somewhere makes it permissible (fully responsible) to pick any of a number of stopping points, none of which is such that we have to stop *there*. But the same sort of difficulty arises as before: it is important that not any old stopping point will do. Surely 'full responsibility' is not demonstrated when (for instance) a scientist continually trusts readings from her complex and sensitive scientific instruments over a period of several years without ever performing any checks on the reliability of the apparatus.

Instead of a notion of *full* responsibility, Wright might try appealing to a notion of *adequate* responsibility. This, however, still needs further comment. One might quite naturally take the standards for adequacy to be in some way sensitive to the context, intentions or interests of the subject and/or the person making the assessment. But anyone tempted by this type of view already has access to more familiar methods of addressing Cartesian scepticism (which moreover seem to promise something more than a merely 'sceptical' solution to the problem: see e.g. DeRose, 1995; Hawthorne, 2004; Lewis, 1996; chapter 4). The appeal to entitlement would then seem to be something of a spare wheel.

There are still things we might say to distinguish this claim from a straightforward claim of strategic entitlement⁹ of a kind that could be accused of establishing nothing more than that ‘non-scepticism is the practical alternative’.

To see what they are, let’s consider (for purposes of contrast) Pascal’s Wager argument that we ought to believe in God, if we want to do well, since believing in God cannot make things substantially worse for us and may make them substantially better. Among the most significant reactions to this type of argument is the following:

O *The kind of justification offered for belief in God is pragmatic rather than epistemic. That is to say, the argument offers pragmatic rather than epistemic reasons for such a belief.*

It might seem as if any response to Cartesian scepticism based on the claim that we have entitlement of cognitive project to trust that S should be subject to an analogous response. Here is an initial attempt to formulate the objection:

*O** *The kind of entitlement proposed in connection with S is pragmatic rather than epistemic. That is to say, Wright offers pragmatic rather than epistemic considerations in favour of such trust.*

The basic thought that *O** tries to capture is that the ground Wright offers for thinking that we are ‘beyond rational reproach’ in trusting in S is that trusting in S is, in some sense, *a good idea*. But Cartesian sceptical argument of the kind Wright hopes to block does not challenge the claim that it is in some sense a good idea to trust that S, only the claim that it is *epistemically* respectable to do so.

There is an obvious sort of reply to objection *O** as it stands. There is an important *disanalogy* between Wright’s argument from entitlement of cognitive project and the simplistic Wager argument cited above, which is that in the former case the various possible consequences of trusting in S are being assessed for their *epistemic* value; non-epistemically valuable consequences such as eternal reward do not come into it. It matters that Wright is talking about entitlement of *cognitive* project, not entitlement of any kind of project whatever. It matters because, on some definitions, it is ‘epistemically rational’ to do X whenever doing X has better epistemic consequences than the alternatives.

‘Epistemic consequentialism’, at a first pass, is the view that the epistemic status of an attitude is determined by the epistemic value of the consequences which result from the subject’s having that attitude.¹⁰ (A recent discussion of this type of position can be found in Percival, 2002.) Consequentialism so characterized could be adopted wholesale (so that the presence or absence of all kinds of epistemic status is determined in this manner), or piecemeal (so that e.g. we are consequentialists when it comes to determining whether a belief is epistemically rational, but not when it comes to determining whether a belief counts as knowledge). I shall only be discussing consequentialism about epistemic rationality.

⁹ Note that here we appeal to the dominance of *trusting* that S, as opposed to the dominance of *acting as if one had a justified belief* that S (cf. p. 183 and see fn. 4 above). This is a necessary first step away from the accusation of having merely established something about practical rationality.

¹⁰ Less roughly speaking, we might prefer to say 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. (Cf. the ethical consequentialist view that the moral status of an action depends on how the value of its consequences compare with the value of the consequences of other available actions.)

If some consequentialist conception of epistemic rationality can be defended, and considerations of cognitive project show that the consequences of accepting *S* are better, epistemically, than those of not doing so, then the response offered above to objection *O** will start to look promising. Even if Wright himself did not intend a consequentialist interpretation of his project, epistemic consequentialism does help to make sense of his thought that entitlement is ‘a type of rational warrant which one does not have to do any specific evidential work to earn’. A consequentialist will be amenable to the idea that the epistemic rationality or warrantedness of an attitude does not (necessarily) depend upon its genesis, or upon its relation to evidence or any cognitive accomplishment of ours; she’ll think that all that matters is the epistemic value of the consequences of having that attitude.

Some caution is needed here, though, since the argument under consideration is not that trusting in *S* *will* have better consequences, epistemically speaking, than not trusting in *S*. Rather, it is that the consequences of trusting in *S* *might be better and cannot be any worse*, epistemically, than the consequences of not trusting in *S*. This is not a serious sticking point, however; we can easily imagine forms of epistemic consequentialism which will embrace the idea that if doing *X* is a dominant strategy with regard to epistemically valuable consequences then *X* is the epistemically rational thing to do.

One other point which needs to be clarified in this context is that Wright’s discussion of entitlement seems to be designed for epistemic internalists (see Pritchard 2005, §3). It might therefore be suggested that any form of epistemic consequentialism which will be of use to us in trying to defend Wright’s claims for entitlement of cognitive project ought to be similarly geared towards internalism. Perhaps the claim should be (e.g.) that if doing *X* is *seen to be* a dominant strategy with regard to epistemically valuable consequences (or would be with the relevant prompting), then *X* is the epistemically rational thing to do.

But in fact, in this context, Wright’s internalist leanings manifest only as a desire to address the problem of *claiming* knowledge where there is dependence on some proposition to which we have only an entitlement; sceptical argument, he acknowledges, does not threaten the claim that we *possess* such knowledge, where knowledge is construed along externalist lines (see his pp. 209–211). Moreover, since he does not appear to want entitlement to be the preserve of the philosopher (pp. 204–205), it can’t be that he imagines that one has to have *seen* (or even be capable of seeing) why one is entitled in order for one to be so entitled.

So for current purposes let us stick with the original (simpler) consequentialist characterization of epistemic rationality, but note that in order to *see* that an acceptance is rational (and hence, presumably, in order to be able to *claim* knowledge in areas where we depend upon that acceptance) we must *see* that doing *X* is a dominant strategy with regard to epistemically valuable consequences. The latter is precisely what Wright is trying to enable us to do.

Are there any card-carrying epistemic consequentialists? Well, Richard Foley (1987, pp. 7–8) argues that it is epistemically rational to do *X* whenever on careful reflection one has reason to believe that doing *X* will promote one’s epistemic goals, provided that all else is equal.¹¹ Although Foley does not use the word ‘consequentialist’ to describe his position, he seems to be committed to an (internalist)

¹¹ For instance, it must not be the case that not doing *X* promotes one’s epistemic goals even more than doing *X* does.

version of consequentialism about epistemic rationality. On Foley's account, an attitude is epistemically rational (provided that all else is equal) to the extent that there is reason to believe on reflection that having that attitude has epistemically valuable consequences, namely the promotion of one's epistemic goals.

One unusual feature of Foley's characterization of epistemic rationality is that it does not immediately focus our attention upon truth, or upon truth-conducive (evidential) reasons for belief, as some do. (See e.g. Kelly, 2003, who takes epistemic rationality to be 'the kind of rationality which one displays when one believes propositions that are strongly supported by one's evidence and refrains from believing propositions that are improbable given one's evidence'.) Foley's focus on truth only appears when we learn that he thinks our only epistemic goal is to believe true propositions now and avoid now believing false ones (p. 8).

However (perhaps surprisingly, given that the novel feature of Wright's proposal is his claim that there is a non-evidential form of warrant), this refusal to focus immediately on truth-conducivity is *not* the feature which makes Foley-style definitions of epistemic rationality particularly significant in this context. What makes them interesting for our purposes is that they do not attempt to cash out the claim that one is epistemically rational in accepting *p* in terms of any claim about one's epistemic relation to *the proposition p in particular*. Rather, on a Foley-style definition, accepting *p* may be epistemically rational just because it promotes our epistemic goals (whatever they are) *in general*, even if it does little or nothing to promote any particular epistemic goals we have with regard to *p* (such as, for instance, the goal of accepting *p* if and only if it is true).¹² Even supposing our sole epistemic goal is to accept true propositions now and avoid now accepting false ones, it appears that we need not have evidence that *p* itself is true in order for acceptance of *p* to be epistemically rational by Foley's lights; we only require that (as far as we can tell on reflection) accepting *p* promotes the goal of now accepting true propositions and not now accepting false ones *in general* (and that other things are equal). It could do that by e.g. resulting in the acceptance of a lot of true propositions other than *p* and the acceptance of a very few false ones, *p* being among the latter.

By Foley's definition, then, Wright can be taken to have shown that accepting *S* is epistemically rational if he has shown that, as far as can be ascertained on reflection, accepting *S* promotes our epistemic goals (and that all else is equal). Other similar definitions will give the same result. I won't concern myself here with whether epistemic value necessarily consists in the 'promotion of one's epistemic goals',¹³ or what those goals are. Instead, let us simply grant for the sake of argument that to do something which is a dominant strategy with regard to the serious pursuit of important cognitive projects (of the kind for which *S* is a presupposition) is to do something which is highly epistemically valuable. (Although it is not obvious that to do something which enables us to carry out these cognitive projects is to do something which promotes Foley's goal of now accepting true propositions and not now accepting false

¹² Indeed, even if it actually *hampers* these goals somewhat, it might still be rational to accept *p* in virtue of the fact that doing so promotes our other epistemic goals to such an extent that this negative effect is outweighed.

¹³ For all I've said so far, it may be that we don't need to characterize epistemic value in terms of goals at all; it may not be a species of instrumental value. Below, however, I will argue for a particular instrumental characterization of the type of epistemic value which I claim is relevant to assessments of epistemic rationality.

ones, there could be other characterizations of epistemic value which will make clear why the needed assumption is true.) Let us even grant that:

D *The way to optimize epistemic value (other things being equal) is to accept S, because that is a dominant strategy with regard to the serious pursuit of certain cognitive projects.*

If anything in this area will help us address objection O*, this will.¹⁴

Suppose, then, that D is granted. The proponent of entitlement must still engage with the worry that, in the cases where we are supposed to have entitlement of cognitive project to accept some proposition, other things are *not equal*. It may be that accepting S has some positive epistemic consequences in that it is a dominant strategy with regard to the serious pursuit of certain cognitive projects. But it may (for all we have yet said) also have negative epistemic consequences, if (e.g.) it results in (or consists in) the acceptance of one or more false propositions. It may moreover be that there are possible situations where *non-acceptance* of S has better epistemic consequences overall, when the epistemic drag factor of accepting false propositions is weighed against the epistemic benefits of optimising the chance of serious pursuit of the relevant projects. If there are such possible circumstances, then although acceptance is a dominant strategy with respect to the serious pursuit of certain cognitive projects, it is not a dominant strategy with respect to optimising epistemic value. (Nikolaj Pedersen raises a point structurally similar to this in an unpublished manuscript.)

But Wright does say that entitlement of cognitive project only gets off the ground in cases where the projects for which S is a presupposition are such that their failure (and presumably their execution where S is in fact false would amount to a failure) would be no worse than their non-execution, and their success would be better. So, assuming (charitably, in order to supply some answer to objection O*) that it is fair to interpret ‘better’ and ‘worse’ here as meaning ‘better and worse as far as epistemic consequences go’, Wright has already specified that entitlement of cognitive project is not present unless it is the case that accepting S cannot make us fare worse with regard to epistemic consequences than non-acceptance, and may enable us to fare better.¹⁵ It may be that in fact there is no proposition like this, in which case we can never have entitlement of cognitive project. But there is something plausible to the

¹⁴ That said, it is in fact far from clear whether acceptance of S is a dominant strategy with regard to the serious pursuit of these projects. Suppose I am a brain in a vat whose senses are being fed information about the world outside the vat in a way that is comprehensible, and helps me learn about the outside world, once I realize that this is what’s going on. Then the best strategy with regard to the serious pursuit of the projects in question (using sensory input to find out about the outside world) is to accept, not S, but the proposition that I am a brain in a vat being fed information in this way. Consideration of this possibility suggests that the strategy of accepting S is not dominant, because there is a possible situation (albeit a somewhat *recherché* one) where some other strategy is preferable to it. However, I am prepared to grant the needed assumption, both because it might be possible to construct other cases where this sort of problem does not arise, and because there are some interesting general issues further down the line that I want to get to.

¹⁵ An interesting question arises here as to whether it could be that S is true, and that therefore the projects which depend upon our accepting S do not (obviously) fail, yet acceptance of S still has negative epistemic value because of the *irrationality* of accepting S without adequate grounds. One might try to respond to this by saying that acceptance of S is *not* irrational because it has good epistemic consequences on balance. But this reply takes the value of the epistemic consequences of accepting S to be settled *in advance* of any considerations of whether acceptance of S is rational. Whereas the suggestion is that the irrationality of accepting S is in fact relevant to the assessment of the value of the consequences.

thought that things are going no worse for us epistemically if we accept S when it is false than if we don't accept it and so never get started on many of our most basic epistemic projects; at any rate, I am happy to allow this assumption in order to see where it leads. I am interested in whether entitlement of cognitive project, supposing we had it, would make our acceptance of a proposition epistemically respectable. I am less interested (for now, anyway) in whether we actually have entitlement of cognitive project with regard to any particular proposition.

My aim, in the rest of this paper, is to argue that *even* if we allow that accepting S is a dominant strategy as far as epistemically valuable consequences go, *and* we are amenable to some form of dominant-strategy-friendly epistemic consequentialism, we might *still* doubt whether we are in a position to give a defence of the epistemic rationality of accepting S.

The basic thought is roughly that we shouldn't be able to establish that accepting S is epistemically rational just by establishing that such acceptance has optimal epistemic consequences *overall*, and without needing to make reference to the *particular* epistemic goals we have with respect to S. The challenge for the next section will be to explicate and defend this intuition.

Proposals for epistemic consequentialists

To stress, my claim is not that epistemic consequentialism is wrong, but that I am suspicious of the *kind* of consequentialist characterization of epistemic rationality that would be needed to vindicate the claims Wright makes for entitlement of cognitive project. We need not, for current purposes, question the idea that the epistemic rationality of an attitude is in some way determined by the consequences of adopting that attitude, as compared with the consequences of not adopting it. I merely suggest that, if we are going to define epistemic rationality in this way, only *some* epistemically valuable consequences should be taken as having the right kind of epistemic value to affect assessments of epistemic rationality. It is not the case that anything which would ordinarily be considered epistemically valuable can be thrown into the equation.

For the sake of analogy, think of an ethical consequentialist—someone initially tempted by simple hedonism but now trying to develop a more sophisticated proposal—arguing that not all kinds of pleasurable consequence are to be taken into account when assessing the moral status of an action. Such a person is still a consequentialist (and perhaps still a hedonist) because she thinks that *only* pleasurable consequences count, although she does not think that *all* pleasurable consequences count. Of course what she needs, in order to motivate the restriction, is some explanation of why some kinds of pleasurable consequence are morally relevant and other kinds are not. Similarly, I will need an explanation of why some kinds of epistemically valuable outcome might be taken to be relevant to epistemic rationality and other kinds should not.

For ease of expression, I shall use the term 'epistemic_r value' to talk about the kinds of epistemic value which are relevant to assessments of epistemic rationality. What I am suggesting is that, even if focussing on all the epistemically valuable consequences of a cognitive act could (for the consequentialist) establish a kind of rationality for that act which we might describe as in *some* sense 'epistemic', this would not be epistemic rationality as we know it. In other words, epistemic_r value is not simply epistemic value, but something more specialized.

A first intuition is that, if we are to be epistemic consequentialists, then, in order to assess the epistemic rationality of an attitude towards a proposition P, we should decide whether adopting that attitude has consequences which are epistemically optimal *in ways which directly concern P itself*. We should not take into account anything which is only epistemically valuable for other reasons, such as consequences whose epistemic value derives from the position we would stand in with regard to propositions *other* than P, or from the position we would be in with regard to cognitive projects other than the project of deciding whether to adopt this attitude to P.

To see the appeal of this, note that we get some very strange results if we allow such extraneous consequences to count for or against the epistemic rationality of believing a proposition. For instance, suppose that some quirky goddess has so arranged things that if I believe P—some proposition which I have no other reason to accept and which is in fact false—then she will arrange for the rest of my life to go so fortunately that all the *other* cognitive acts I ever perform will be absolutely brimming over with all the features that generate epistemic value (whatever they are). However epistemically valuable this consequence of believing P might be, and even if I knew all about the goddess's intentions, the acceptance of P would still be epistemically irrational.¹⁶ If we do not restrict the notion of epistemic_r value in something like the way I've suggested, this intuitive claim will be false: epistemic rationality will drop out of the high epistemic value of the consequences of believing P.

But once we discount the kinds of epistemic value which I have been classifying as 'extraneous', what will be left? A defender of Wright may be concerned that, by my lights, nothing will count as the right kind of epistemically valuable outcome of accepting P except the fact of having *thereby accepted a true proposition* (or, more plausibly: one which is likely to be true given our epistemic situation). If that's right, then it might seem that I am begging the question against Wright here, by assuming that a defence of the epistemic rationality of accepting S must make some mention of S's truth (or probable truth given our epistemic situation). That looks dangerously close to simply assuming a negative answer to the interesting question he is exploring, namely, whether there could be non-evidential warrants.

I shall try to address this worry about question-begging by attempting to locate independent grounds for the intuition that the kinds of consequence I've described as extraneous really are such. Instead of assuming that those consequences are extraneous because they do not relate to the (probable) truth of the proposition under consideration, I shall argue that:

- C: *If we are epistemic consequentialists, we ought to think that the epistemic_r value of a cognitive act¹⁷ depends upon its promotion of those aims which it has in virtue of its being the kind of cognitive act it is.*

¹⁶ Foley's restriction of the goal of belief to *now* believing true propositions and not *now* believing false ones is supposed to help with this sort of problem (see Foley, 1987, p. 8). But I don't see how it does; the same issue arises if we have the goddess agree somehow to supply us with masses of true and highly epistemically valuable beliefs *simultaneously* with our acceptance of P.

¹⁷ I follow Percival (2002) in using 'cognitive act' in a way that doesn't imply that such acts are optional or voluntary (see his p. 122). A cognitive act, for my purposes, is anything which it makes sense to assess for epistemic rationality (e.g. belief that P, disbelief of P, suspension of belief as to whether P, trust that P). From this it follows that epistemic_r value can only accrue to cognitive acts, although (for all I say here) any kind of act might be *epistemically* valuable.

If C can be sufficiently motivated, we can resist the thought that Wright has provided a defence of the epistemic rationality of S, without assuming that every such defence will be evidential, and hence without begging any questions. C will enable us to argue that, even if accepting S is a dominant strategy with regard to epistemic value in general, it is not a dominant strategy with regard to epistemic_r value. For (as I'll argue), in the case of acts of acceptance, the aims referred to in C relate to the specific proposition accepted, and not to cognitive projects other than that of deciding whether or not to accept that proposition.

Why should the aims relevant to epistemic_r value be the ones that cognitive acts have in virtue of being cognitive acts of a certain kind, as C says? Even if we accept that *some* clarification is needed of the consequentialist claim that the epistemic status of an attitude is determined by the epistemic value of its consequences, surely there are various options that might be considered.¹⁸ One might, for instance, propose that:

C': If we are epistemic consequentialists, we ought to think that the epistemic_r value of a cognitive act depends upon its promotion of those aims which it has in virtue of being a cognitive act.¹⁹

Or:

C'': If we are epistemic consequentialists, we ought to think that the epistemic_r value of a cognitive act depends upon whether it is epistemically optimal in ways that directly concern the proposition under consideration.²⁰

And no doubt many other ingenious proposals could be imagined.

I don't have a general argument to offer here in favour of C (nor is it possible to discuss the alternatives in any detail without too much of a diversion from my main theme); but let me mention an analogy which I find motivating. The analogy is between acting in ways which are epistemically rational (thinking well) and acting in ways which are rational in the context of playing a game of chess (playing well). Suppose we are interested in assessing how good my last chess move was, in order to see whether I am playing well. Suppose also that we are consequentialists about good play in chess: we think that a move is good insofar as it has good consequences. Obviously, some kinds of good consequences which my last move might have had are not supposed to count for this purpose. For instance, if I moved my queen and this made one of my spectators very happy because he thinks the queen's move is aesthetically pleasing, that would be a good consequence, but not one relevant to an assessment of my standard of play. Which consequences *are* relevant? I suggest: the ones which relate to those aims which the moves have in virtue of being moves in a game of chess. The aim which moves in a game of chess have in virtue of being such

¹⁸ Thanks to a second referee for encouraging me to clarify this point.

¹⁹ See my discussion of Percival's (V) on p. 35 below.

²⁰ Note that, although I think *C''* is true, I do not think it is maximally elucidatory. For one thing, C helps explain the truth of *C''*. For another, *C''* allows for (though it does not entail) the possibility that an act's rationality can be affected by consequences concerning the proposition under consideration which have nothing to do with promotion of the aims the act has in virtue of being the kind of cognitive act it is—a scenario I think we should rule out, for reasons to be explained shortly.

as to put the player making them in a better position²¹ from which to win the game. So, for the purposes of assessing whether my move was a good move, the relevant consequences are those which in some way promote or hinder this aim.

By way of additional explanation of my choice of C over the various alternatives, I should note also that the idea that a defence of the epistemic rationality of a cognitive act must make some appeal to what is essential to the cognitive act in question is not a new one, but one that has also appealed to others who were not concerned with the particular problem I have set myself. Percival (2002) has considered something in the vicinity of my proposal within a consequentialist framework. (I shall turn to his discussion in a moment.) And Hieronymi (2005) has suggested more broadly that the distinction between ‘the right kind of reasons’ and ‘the wrong kind of reasons’ for belief is that responsiveness to the right kind of reasons is what is *constitutive* of believing.²²

Percival (p. 146) suggests that:

- (V) The concept of [epistemic]²³ value is the concept of that (sort of outcome) for which cognition aims.

While I think this may be heading in the right direction as far as the characterization of *epistemic* value goes, we will need to refine it in order to understand *epistemic_r* value. For one thing, different cognitive acts (e.g. believing that P and believing that Q, or believing that P and withholding belief from Q) may have different aims, and the epistemic rationality of a particular act would seem to be tied to the promotion of the *particular* aims which *that act* has essentially.

To motivate this last thought, consider another analogy: that between epistemic rationality (thinking well) and rationality in the context of completing a crossword puzzle (solving well). Acts of entering words into the grid seem to have the common essential aim of making it the case that I have entered as many right answers as possible, but this is true in virtue of the fact that any particular act of entering a word into a particular position in the grid has the essential aim of making it the case that I have entered the right answer *in that position*. And it is the latter aim promotion of which seems to characterize good crossword-solving. If I enter something which I have no reason to think is the right answer, just because I know that if I do that my friend the quirky goddess will enable me to work out the right answers to all the other clues, that does not make for a good bit of crossword-solving. It wasn’t rational (in the intended sense) to enter that answer, because I didn’t have any reason to think it was correct. Success in the relevant dimension of assessment is measured by the promotion of the particular essential aim of entering a correct answer *in the location under consideration*.

²¹ A better position, that is, as far as the rules and strategy of the game itself go. Obviously some moves could annoy my opponent, making him more likely to stop playing, so that I cannot win. But it is not the case that chess moves have the aim, in virtue of being moves in a game of chess, of helping me keep my opponent on an even keel temperamentally.

²² Although this suggestion is interestingly related to the present discussion, I don’t think it can work. Belief, in a sufficiently irrational person, could be responsive to anything at all. But if we specify that belief is constituted by *rational* responsiveness to the right kind of reason, we seem just to have pushed the question back. We will now owe a definition of ‘rationality’ which distinguishes the epistemic (‘right’) kind from other (‘wrong’) kinds of rationality.

²³ Percival says ‘cognitive’ here, but for current purposes I take that to mean the same as what I mean by ‘epistemic’.

I propose that we refine (V) so as to make reference to the essential aims of particular cognitive acts. (I also propose to eschew mention of concepts here, since this raises a set of difficult but tangential issues.) As a first pass, then, let's consider:

- (VP) The epistemic_r value of a *particular act* of cognition is measured by the achievement of that (sort of outcome) for which that act of cognition aims.^{24,25}

(VP) still needs careful interpretation, however, because any cognitive act can have various aims, not all of which are such that their promotion would be relevant to epistemic_r (or even epistemic) value. (Percival notes something similar with respect to (V), on p. 146.) Perhaps one of my aims in accepting a proposition P is to make myself happier, but the achievement of this aim should not count in the balance when we are assessing the epistemic_r value (or even the epistemic value) of accepting P.

Which aims are we interested in, then? Well, for reasons I have already described, I am inclined to agree with Percival that the hope of epistemic consequentialism 'must be . . . that a unique and objective aim of cognitive acts is given by their nature qua cognitive acts' (p. 147). The aims which are of interest to us are the aims which cognitive acts have in virtue of being those cognitive acts.

It is not clear, however, whether Percival thinks the hope of consequentialism must be that cognitive acts in general have some one aim given by their nature as cognitive acts, or whether the hope is that each kind of cognitive act has some one aim given by its nature as a cognitive act of that kind. But the latter is surely enough for consequentialist's purpose; there is no obvious reason to impose the more stringent constraints implied by the first reading. So I propose that we refine (VP) to give:

- (VP*) The epistemic_r value of a particular act of cognition is measured by the achievement of that (sort of outcome) for which that act of cognition aims in virtue of being a cognitive act of the relevant kind.²⁶

It is worth noting here that (VP*) should serve to address a worry raised by Kelly (2003) concerning attempts to define epistemic rationality in terms of the goals or

²⁴ It is not clear to me whether Percival thinks there is a notion of the kind of value an individual act of cognition has in virtue of succeeding in its aim, as well as a notion of epistemic or cognitive value in general. As well as giving the above general account (V) of cognitive value, he also states that 'the epistemic consequentialist's ranking of consequences according to their cognitive value is a measure of the degree to which *the cognitions in question* succeed in achieving *their aim* (p. 146, emphases added). These two notions of cognitive value seem potentially very different to me, but Percival does not explain how they relate to each other. Maybe he is assuming that all cognitions have the same aim. But even if they all have *similar* aims—say, true belief in the target proposition—the aims are plausibly different insofar as the target propositions are different.

²⁵ Let me note in passing here that, contrary to initial appearances, this sort of approach need not create a tension with holism about epistemic rationality (although some adjustments would be needed if one favoured holistic approaches). Holists can take the 'particular acts of cognition' whose epistemic value they are assessing to be acts of believing (or accepting, or whatever) whole theories, rather than propositions which form part of a theory. They could then allow that, derivatively, belief (for example) in a proposition counts as rational if that proposition is part of a theory that it is rational to believe.

²⁶ I have not here attempted a consequentialist characterization of epistemic value, but I am inclined to think it might look something like:

- (V*) Epistemic value is the measured by the promotion of that (sort of outcome) for which cognition aims in virtue of being cognition.

aims of cognitive acts. Kelly is concerned that whatever epistemic goals we appeal to in our definition, we run the risk that some individual could simply lack those goals, and hence fail to be characterizable as rational or irrational. The objection is that even if such an individual has beliefs which we clearly want to categorize as irrational, we cannot do so, because he does not have the epistemic goals that give rise to epistemic rationality and irrationality. This doesn't seem right: one ought not to be able to duck out of being accused of epistemic irrationality just by lacking a certain kind of goal.

The response I favour is that, if the relevant aims are aims which cognitive acts are *bound* to have, just in virtue of being those sorts of acts, then anyone who lacked the relevant aims simply would not be performing the corresponding acts.²⁷ There is something quite intuitive about this. Suppose that an essential aim of the act of adopting a belief is that it constitute an act of adopting a belief in a proposition which is probably true given one's epistemic situation. On the view under consideration, if whatever cognitive acts a subject A performs with regard to proposition P do not have this aim, then they simply cannot be acts of belief. To believe is to aim at the truth. If A does not aim at the truth, A does not believe.

Nevertheless, Kelly (p. 631) attempts to resist this sort of response to his worry, suggesting that the argument:

I have beliefs about x
The aim of any belief is truth
Therefore I have the aim of having beliefs about x which are true

is no better than:

I have a heart
The aim of any heart is to pump blood
Therefore I have the aim of having a heart which pumps blood

where the latter is clearly fallacious (as I could be trying to commit suicide by stopping my heart from pumping blood).

The way to defuse this reply, I think, is to note that the second premise of the original argument is intended to be equivalent to:

To have a belief about x is to have the aim of believing truly about x.

When we talk about 'the aims of cognitive acts', we must remember that these acts (voluntary or not) are acts *of ours*, and that therefore their aims (unlike the aims of parts of our bodies) are, automatically as it were, our own aims. If I am right to think that any plausible version of consequentialism about epistemic rationality must employ a notion of *relevant* epistemic value, epistemic_r value, as given by something

²⁷ A question which remains, however, is whether the *non*-performance of an epistemic act could be irrational. (Plausibly it could. For instance, faced with obviously overwhelming evidence for P and no countervailing evidence, it seems irrational not to believe P.) Since I am only attempting to give a consequentialist characterization of epistemic rationality *for cognitive acts*, it might seem that this isn't directly of concern. However, it does raise relevant issues. For in order to assess the epistemic rationality of a cognitive act we are supposed to be assessing its epistemic_r value compared to that of the alternatives. If a relevant alternative is non-performance of the act, the epistemic_r value of non-performance must be assessable somehow. And there will need to be some way of weighing the epistemic_r value of non-performance against that of performance. I am, however, prepared to grant for the sake of argument that some such weighting can be achieved, and even that if Wright has succeeded in showing that accepting S is epistemically_r valuable, we can take it as read that the epistemic_r value of non-acceptance will be lower. My question is purely whether we have been given any reason to think that acceptance is epistemically_r valuable.

like (VP*),²⁸ then, I think, no version of it will enable us to count what Wright says in defence of our acceptance of S as sufficient for a defence of its epistemic rationality. For all Wright shows is that accepting S is a dominant strategy with regard to the serious pursuit of important cognitive projects. But *even if we allow* that accepting S is therefore optimally epistemically valuable, Wright may not have shown that acceptance is optimally epistemically_r valuable, i.e. optimally epistemically valuable in ways that are relevant to an assessment of the epistemic rationality of accepting S. He may not have said anything relevant to an assessment of the epistemic_r value of accepting S at all. For he gives us no reason to think that accepting S will promote the aims which that act of acceptance has in virtue of being an act of acceptance.

But perhaps such an argument could be given; after all, I haven't yet said anything about which aims an act of acceptance of S has in virtue of being an act of acceptance. In fact, however, I think those aims are such as to make it implausible to suppose that Wright has said anything relevant to the epistemic_r value of accepting S. For one thing, despite the differences Wright claims between acceptance and belief (see his section II), when acceptance is characterized as *trust*, as in the discussion of entitlement of cognitive project, it will seem that the essential aim of any act of acceptance of P is to thereby accept something which is probably true given our epistemic situation, and Wright explicitly tells us nothing relevant to the question of whether the act of accepting S succeeds in this respect.²⁹ I shall offer some thoughts in support of this claim about the essential aim of acts of acceptance in a moment. But note that for current purposes it is sufficient to establish the weaker claim that [being a dominant strategy with respect to enabling the serious pursuit of certain cognitive projects] is *not* among the essential aims of acts of acceptance. This is enough to show that entitlement of cognitive project isn't the right sort of thing to underwrite a defence of the epistemic rationality of accepting S. To see that the point holds, it is enough to note that some acts of acceptance do not have any aim of this sort. For example, some people trust (without evidence, but in a way that precludes doubt) in the existence of God. Yet it would be no reason to think that trust misplaced if it transpired that trust in God was not a dominant strategy with respect to enabling the serious pursuit of certain cognitive projects. This act of acceptance does not have this sort of aim.

Perhaps it might be proposed that acceptance of those particular propositions to which we have entitlement of cognitive project is a *special* cognitive act which *does*, by its very nature, aim simply at enabling the serious pursuit of important cognitive

²⁸ For completeness's sake, let me add that I would also imagine that this notion of relevant epistemic value is the one the consequentialist should use to define various other possibly non-equivalent terms of epistemic approval, such as 'epistemically justified'. It would be an option, I suppose, to say that X is epistemically *justified* if it is optimally epistemically valuable, although X is only epistemically *rational* if it is optimally epistemically_r valuable. But I don't think that this use of terminology sits comfortably with our pretheoretic notion of epistemic justification.

²⁹ The reader might wonder whether I am here simply ignoring Wright's claim (supposed to be partly definitive of the notion of acceptance under consideration) that acceptance is not 'essentially rationally controlled by evidence' (p. 183). I am not; rather, I am attempting to undermine his claim to have made it plausible that there is any kind of cognitive act which involves trust, excludes doubt *and* is not rationally controlled by evidence (broadly construed, as described on pp. 174–175). It is not possible to *define* this kind of act into existence; so if all three of the characteristics just listed are really supposed to be definitive of acceptance, my point will be that Wright is attempting to define acts of acceptance as having an inconsistent set of properties. Which will simply mean that there are no acts of acceptance.

projects. I can't imagine anyone finding this line of response particularly appealing, but let's see where it leads.

I am inclined to reply that no act of this sort would be properly characterized as acceptance in the sense Wright is interested in. But that inclination is due to my belief that acceptance in this sense essentially aims at being an acceptance of something which is probably true given one's epistemic situation. So let me now attempt to defend that claim (abbreviating it, for the sake of clarity, with the slogan 'acceptance essentially aims at probable truth').

The first thing to note is that if acceptance does not aim at *something* like probable truth, then it is hard to see why anything other than a merely pragmatic defence of accepting S should be desirable (where a merely pragmatic defence is the sort that could be generated by saying that those who accept S are rewarded with chocolate). Yet all parties to the current debate agree that it is important to be able to show that we can give something more than a merely pragmatic defence of our accepting propositions like S.

The fact that acceptance of the kind under consideration here necessarily involves trust (p. 194) and excludes doubt (p. 193) also goes some way to suggesting that in accepting a proposition we do in fact aim to accept something which, as far as we can tell, is true. To trust that S is to trust that S is *true*, and such trust is *misplaced* if S turns out not to be true. Were it not the case that we aim to trust only what is probably true given our epistemic situation, there would be nothing misplaced about trusting in a proposition that has turned out to be false; but clearly there is something wrong with that. Similarly, adopting an attitude towards S which precludes doubt about S is to preclude doubt that S is *true*, and is misplaced if S turns out to be false. But were it not the case that we aim to preclude doubt only about what is probably true given our epistemic situation, there would be nothing misplaced about precluding doubt about a proposition which has turned out to be false.³⁰

It is a further question whether probable truth is one of the aims that an act of acceptance has *in virtue of being* an act of acceptance. But if the involvement of trust and the exclusion of doubt are of the essence of acceptance, and this is sufficient to make acceptance an attitude which aims at probable truth, then it does seem as if acts of acceptance aim at probable truth by their very nature. So any special kind of cognitive act which by its nature aims solely at the enabling the serious pursuit of certain cognitive projects is not an act of acceptance.

A more subtle proposal which we might consider is that, by its nature, the special act of accepting one of these fundamental propositions aims *both* at probable truth *and* at enabling the serious pursuit of certain cognitive projects (where these aims are so weighted that achieving the latter aim is sufficiently epistemically_r valuable to make acceptance more epistemically_r valuable than non-acceptance, regardless of whether the former aim—probable truth—is achieved or not). But this sounds really ad hoc; I'm not aware that I perform any cognitive acts like this and find it hard to see any motivation for thinking that I do other than that generated by a desire to rescue the claim that we have entitlement of cognitive project in respect of these propositions. In addition, I doubt intuitively whether any such act would be properly characterized as

³⁰ Wright's clause (i)—see Introduction above—shows that he believes that acceptance cannot be epistemically rational when we have sufficient evidence to believe that S is false. I am trying to push the issue a bit further, and ask for an explanation of *why that is*. My proposed answer is that it is because acceptance essentially aims at probable truth, and epistemic rationality, however we characterize it, is sensitive to the essential aims of cognitive acts.

an act of acceptance in the sense Wright intends. Acceptance (trusting beyond doubt) does not seem, by its very nature, to have any aim *other* than probable truth. At the very least, the ball would seem to be firmly in Wright's court to make it plausible that there are cognitive acts of the kind required.

Conclusion

I can now explain my objection to Wright's claim for entitlement of cognitive project, without begging any questions by assuming at the outset that any genuine defence of the epistemic rationality of accepting S will have to be evidential. I note instead that, even granting that we are prepared to accept some form of epistemic consequentialism, our being 'entitled' to accept S in the way Wright envisages tells us nothing about whether accepting S has optimal epistemic value *in the ways that are relevant for assessments of epistemic rationality*, because it tells us nothing about whether accepting S will promote the essential aim(s) of accepting S. And for that reason the entitlement story does not make for a defence of the epistemic rationality of accepting S.

Of course, some part of the question of whether to call what Wright offers a defence of *epistemic* rationality (rather than some other kind of rationality) is terminological; we could decide to use 'epistemic rationality' in such a way that Wright's defence counts as a defence of epistemic rationality because it focuses on the positive *epistemic* consequences of acceptance. However, the suspicion that it is not properly characterized as a defence of epistemic rationality (in the sense we all know and love) is confirmed when we consider whether it enables us to resist the challenge to epistemic rationality raised by Cartesian sceptical argument. As far as I can see it does not.

According to Wright, sceptical worries are 'tamed' by the thought that, even if no evidence for S is forthcoming, I can still be 'rationally entitled' to take S 'on trust' when engaged in 'any broadly empirical cognitive project' (p. 195). But which worry exactly is being tamed here? Not the worry that we cannot give reasons for accepting S which make us confident that S is true, for, as he acknowledges, that worry remains. (That is why his solution is a 'sceptical' one.)

Perhaps the 'tamed' worry is supposed to be a worry about whether we can show that accepting S is preferable to non-acceptance in terms of overall epistemic consequences. But this is not a familiar worry; the problems raised by sceptical argument of the kind which encourages us to doubt whether S is true certainly don't present themselves as being problems of this nature.

Wright has suggested in discussion that the tamed worry is a worry about whether we can live with the fact that the cornerstones of enquiry are always beyond the reach of enquiry, and 'still regard ourselves as rational'. But as I emphasized at the beginning of this paper, a good deal hangs on what *kind* of rationality is at issue here. The worry posed by Cartesian scepticism seems to be a worry about epistemic rationality, and my purpose has been to argue that it's hard to see how that worry is tamed by appeal to the notion of entitlement of cognitive project.

Acknowledgements I am grateful to all who participated in various research seminars on this topic at the University of St. Andrews while I was preparing this material, particularly Daniel Nolan, Nikolaj Pedersen, Duncan Pritchard and Robbie Williams, and to an audience at the Joint Session of the Mind Association and the Aristotelian Society 2005. Ideas and suggestions arising from these

sessions and subsequent discussions inform many parts of this paper. Particular thanks to Crispin Wright for comments on an earlier draft. Comments from two anonymous *Synthese* referees have also been helpful.

References

- DeRose, K. (1995). Solving the skeptical problem. *The philosophical review* 104; repr. In Sosa, E. & Kim, J. (Eds.), *Epistemology: An anthology*. Oxford: Blackwell, pp. 482–502.
- Field, H. (2000). Apriority as an evaluative notion. In Boghossian, P. & Peacocke, C. (Eds.), *New essays on the a priori*. Oxford: Clarendon Press, pp. 117–149.
- Foley, R. (1987). *The theory of epistemic rationality*. Cambridge, MA: Harvard University Press.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Clarendon Press.
- Hieronymi, P. (2005). The Wrong Kind of Reason. *The Journal of Philosophy* 102, pp. 437–457.
- Kelly, T. (2003). Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research*, LXXVI, pp. 612–640.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy* 74; repr. in his *Papers on metaphysics and epistemology*, 1999, Cambridge University Press, pp. 418–445.
- Percival, P. (2002). Epistemic consequentialism I. *Aristotelian Society Supplementary Volume LXXXVI*, pp. 121–151.
- Pritchard, D. (2005). Wittgenstein's *On Certainty* and contemporary anti-scepticism. In Moyal-Sharrock, D. & Brenner, W. H. (Eds.), *Readings of Wittgenstein's On Certainty*. London: Palgrave Macmillan, pp. 189–224.
- Reichenbach, H. (1938). *Experience and prediction*. Chicago: University of Chicago Press.
- Wittgenstein, L. (1969). *On certainty*. Oxford: Basil Blackwell.
- Wright, C. (2004). On epistemic entitlement: Warrant for nothing (and foundations for free?). *Aristotelian Society Supplementary Vol. LXXVIII*, pp. 167–212.