

Epistemic Norms and Natural Facts *

C. S. Jenkins

Forthcoming in *American Philosophical Quarterly*

Abstract

Has enough been done to ‘naturalize’ epistemic normativity if, for every true epistemically normative claim, the fact which it picks out and which makes it true has been identified, and that fact is shown to be a natural fact? It's tempting to say 'yes'. This sort of view appears to raise an analogue of the Open Question Argument faced by ethical naturalists. (Similar worries have also been raised about naturalist views on other topics.) But one can respond, as is familiar from these other debates, that it need not be part of the view that any epistemically normative claim *means the same* as a claim which more explicitly picks out the relevant natural fact, and that the identity between the facts which make these two claims true need not be something we can learn by conceptual examination, or something which is at all obvious. The paper compares this form of naturalism about epistemic norms with other available forms. It does not defend any particular proposal as to which natural facts make true epistemically normative claims, but discusses the merits of the strategy in general terms.

I The Strategy

Naturalism, for the purposes of this paper, is the position which holds that some scientific world view is approximately correct, so that there exist no supernatural or otherwise spooky entities, properties, events or other phenomena. If naturalism is correct, then the world is the way that best science says it is (where best science need not be identical to any actual science, whether present or future), and there is nothing more to be truly said about the world.

Normativity is sometimes felt to present a challenge to naturalists: normative facts seem to be spooky or non-natural, in a way which non-normative facts are not. Moreover, normative facts do not appear to belong to the subject matter of science, which seems to be a purely descriptive enterprise: science does not, at least *prima facie*, deal in facts about how things *should* be, it just tells us how things are. The normative facts, if any there are, seem to belong to a different realm.

Normative discourses are many and various, but this paper concentrates on epistemic normativity: the kind of normativity we seem to be attributing when we say that - regardless of any consideration of moral status or practical consequences - we *ought* to fit our beliefs to our evidence, justified beliefs are *in good standing*, blind trust in unreliable informants is *wrong*, it is *irrational* to believe an explicit contradiction, and so on. It is hard to deny the existence of at some clear cases of true epistemically normative claims. If these claims are made true by epistemically normative facts, are these the kind of spooky facts that undermine naturalism? The point of this paper will be to present and discuss a certain strategy for construing epistemic normativity in a naturalistically acceptable way, and point out some of its advantages relative to other naturalizing strategies. There is nothing new about the strategy itself: it is familiar from

discussions of ethical normativity, and similar moves have been made in other debates. What is new is the application of this strategy to epistemic normativity.

In ethics, one relatively popular strategy for naturalizing normativity (rendering it naturalistically acceptable) is to argue that true morally normative claims pick out and are made true by natural facts, such as facts about the maximization of some value which can be identified with some naturalistically acceptable property. For instance, one might say that facts about the maximization of utility are the facts which make true morally normative claims. If this strategy succeeds, the only facts we need to posit in order to account for the truth of morally normative claims are naturalistically respectable facts.

An analogue of this strategy looks promising in connection with claims of epistemic normativity. The best-known strategy for rendering epistemic normativity naturalistically acceptable involves claiming that it is merely a species of instrumental normativity. (This view is defended by Quine, Laudan and others.) It will be argued here that the envisaged strategy is importantly different, and that it overcomes certain objections to this more traditional approach. It will also be distinguished from other – extant or possible – epistemic naturalizing strategies, and some of its advantages over them will be noted.

The naturalizing strategy to be discussed here rests on making some significant association between epistemically normative claims and certain natural facts. A preliminary question to be addressed, therefore, is: what kind of association, exactly? To illustrate with a familiar ethical example, let's suppose that the claim:

O1: *X ought to help Y*

is said to be made true by the same fact as makes true the claim:

M1: *It would maximize overall utility if X were to help Y.*

What does this proposal amount to? Does it claim that the sentence we used in order to express O1 means the same as that used to express M1? That O1 and M1 pick out and are made true by the same fact or worldly state of affairs? Is it that O1 is true in virtue of the truth of M1? Or that M1 explains why O1 is true? Maybe the idea is just that O1 is true *iff* M1 is true? Or that the ought facts supervene on the utility facts? For current purposes, the ought facts supervene on the utility facts just in case any worlds which differ in ought facts also differ in utility facts. Kim (1988) suggests that epistemology can count as naturalized provided we have established that the epistemically normative facts supervene on the natural facts in something like this sense.

It is fairly clear that the last two proposals – the biconditional proposal and the supervenience proposal – are too weak, given that the aim is to show why the truth of O1 is naturalistically unproblematic. The fact picked out by O1 might still be spooky for all they say. For all the supervenience thesis tells us, the epistemically normative facts are spooky facts which supervene on the natural facts. And the biconditional could be true just because the spooky fact which makes true O1 obtains if and only if the natural fact which makes true M1 obtains. The antepenultimate, explanatory, proposal is also too weak: for all it says, it could be that the citing the natural fact M1 is a way of explaining why the spooky fact which makes O1 true obtains.

On the other hand, it does not seem necessary to go so far as to claim that the two sentences used to express claims O1 and M1 mean exactly the same:¹ it would be enough if these two sentences had different senses, but at the level of reference picked out the

same worldly fact, and were both made true by that fact. Versions of this kind of view, which (sometimes together with other doctrines) have attracted such labels as 'Cornell realism' and 'synthetic moral naturalism', are defended by Sturgeon (1985), Railton (1986), Boyd (1988), Brink (1989) and others.

The locution 'in virtue of' will be avoided in the remainder of this paper, since it tempts one to blur the distinctions between the various options described here. For instance, both the view that M1 explains why O1 is true and the view that M1 and O1 are made true by the same fact seem describable as views according to which O1 is true 'in virtue of' the truth of M1.

If one takes the moderate route of saying that O1 and M1 pick out the same fact despite having different senses, one can move to address the Open Question Argument (whose *locus classicus* is Moore 1903): that is, the objection that a natural fact like that which makes true M1 does not adequately capture the normativity of O1, because even when it is settled that X's helping Y would maximize overall utility it remains an open question whether helping Y is what X ought to do. This objection can be addressed by saying firstly that, if all that is meant is that it is not *conceptually* and/or *obviously* trueⁱⁱ that the truth of O1 is settled once the truth of M1 is settled, then this is no objection to the thesis that O1 and M1 are made true by the same fact. For this thesis doesn't imply that the sentence used to express O1 means the same as that used to express M1, and it need not be obvious that two sentences with different senses pick out the same fact even when they do. But if there is supposed to be more to the objection than this, it is hard to see what it is supposed to be. To say that the question of O1's truth remains *genuinely* open when M1's truth is settled is to beg the question against the view that the same fact

makes true these two claims, not to raise an objection to that view.

Note that the proposed association between M1 and O1 is asymmetric in that we are supposed already to have an idea of what kind of fact makes true M1, and we are to use this to guide our conception of what kind of fact makes O1 true. The mere claim that both O1 and M1 are made true by the same fact does not settle things this way, since one might conclude on that basis, not that O1 is made true by the same *natural* fact as M1, but that M1 is made true by the same *spooky* fact as O1. Clearly, it is the former conclusion that the naturalist intends to defend.

Those who see no mileage in trying to naturalize moral normativity by associating morally normative claims with natural facts will no doubt be similarly suspicious of this type of strategy as applied to epistemic normativity. Therefore, what follows will be of most interest to those who have some sympathy (or at least, are not out of sympathy) with this strategy in the moral case. Nevertheless, it should be of interest to everyone to think about how the issues play out in this arena, and to note the interactions with other positions which are being discussed in the epistemic case.

A brief note is in order here concerning the distinction between *naturalizing* normativity and *getting rid* of it. The strategy to be defended here is supposed to be a way of maintaining that epistemic normativity is a genuine, distinctive and non-instrumental kind of normativity, that claims of epistemic normativity are true, and that they are made true by facts in the world. The idea is that they are made true by the same facts as certain other claims which are more obviously naturalistically acceptable. This does not amount to showing that there are no epistemically normative facts. On the view under consideration, the epistemically normative facts are not eliminated but simply

identified as certain natural facts.

For those who like the terminology of reductionism, the view might be described as one on which epistemically normative facts are 'reduced' to natural facts. Some caution is necessary with this terminology however, since there is not supposed to be any '*conceptual* reduction'; that is to say, it is not assumed that epistemically normative concepts are the same as the relevant naturalistic concepts, nor that the relationship between epistemically normative claims and the corresponding natural-sounding claims can be discovered just by thinking about about the concepts involved. (More on this and related issues in §IV below.) As a reductionist strategy, it contrasts with strategies which claim merely that epistemically normative facts supervene on natural facts. But it will presumably *imply* supervenience: normative and natural facts are identified on this strategy, and it is presumably trivial that facts supervene upon themselves.

II Some Versions

Consider the following claim:

I1: We ought to avoid believing any inconsistent proposition.

There are many things that could be going on if this claim is true. It might be, for instance, that we ought not to believe inconsistent propositions because if we avoid doing so we will win £1000, or because to do so would be somehow immoral. Set this sort of thing to one side, however, and focus instead on the reading (or readings) of I1 on which the normativity involved is distinctively epistemic.

In this section, two kinds of naturalistically respectable fact will be discussed which might be suggested as candidates for the facts which epistemically normative

claims like I1 pick out and which make such claims true. But it should be noted that for those who do not think either of these specific proposals is correct, various other options are available, compatible with the overall strategy. One simply has to select a different kind of natural fact to associate with claims of epistemic normativity. Those cited here are merely examples, included to help make the overall strategy clearer.

Of course, there are many different kinds of epistemically normative claim: 'ought' claims and 'should' claims, claims about what it is 'right' to believe, claims about what it is 'rational' to believe, claims about what it is 'permissible' to believe', and so on. It might be that different kinds of natural facts are associated with the different kinds of epistemically normative claims on this list.

Equally importantly, a single normative sentence might have two (or more) distinct acceptable readings, both (all) of which are distinctively epistemic. The sentence used above to express I1, for instance, might be taken to have both 'internalist' and 'externalist' readings, and these might be taken to diverge in truth-value.ⁱⁱⁱ For current purposes, let's take internalism to be the view that one can tell 'from the inside' what one epistemically ought to do (or what is right, or rational, or ...), and externalism to be the view that one cannot necessarily tell from the inside what one epistemically ought to do (or what is right, or rational, or ...). Distinctions of degree can also usefully be drawn between various positions which are not fully internalist in the above sense. Proposals will be described here as 'more internalist' or 'less internalist' depending on the *extent* to which, according to them, it is possible to tell from the inside what one epistemically ought to do.

It is far too broad a question to settle here whether all our epistemically normative

talk and thought deals in internalist notions, or whether it deals exclusively in externalist notions, or whether some of our normative notions are internalist and others externalist. (The present author leans towards this last view.^{iv}) For now, note only that someone interested in defending the overall strategy outlined in this paper is free to adopt any of these three positions, providing she can find suitable natural facts to serve as the required normative facts. Of the two suggestions to be made in this section, one is clearly externalist in flavour, while the other might be made more externalist or more internalist, depending on how one construes its appeals to evidence and to probability.

The first suggestion is motivated by the fact that the plausibility of I1 clearly has *something* to do with the fact that inconsistent propositions cannot be true. One attractively simple naturalistic proposal is to say that I1 is made true by the same fact as: N1: An inconsistent proposition is not true.

This doesn't rescue naturalism all by itself. For we'll need to be confident that there is a naturalistically acceptable way to understand the notion of truth in N1. This is not the place to address this issue, however. The same sort of point arises in the ethical case in connection with counterfactuals like M1.

This association between I1 and N1 could then be taken as a particular case of a general principle that claims of the form:

G1: S's belief that p is in good standing epistemically (or would be if S had such a belief)

are made true by the same facts as claims of the form:

T1: p is true.

Similarly, we might think that claims of the form:

B1: S's belief that p lacks epistemic good standing (or would if S had such a belief)

are made true by same facts as claims of the form:

T2: p is not true.

Since inconsistent propositions are not true, these general principles would imply that beliefs in inconsistent propositions are not in good standing, so that one epistemically ought not to believe any inconsistent proposition. These principles make sense of the intuitive idea that the 'aim', 'norm' or 'goal' of belief is truth, and that the way for beliefs to be 'wrong' is for them to be false.

Many will no doubt feel, however, that this association between normative claims and natural facts is too simplistic. They will argue that excellent but misleading evidence for p can make a belief in p fully epistemically correct, even though p is in fact false. For example, suppose Nick has every reason to trust Gale when she tells him that her socks are blue, and he has no other evidence on the subject. Nick comes to believe on the basis of Gale's testimony that her socks are blue, when in fact they are green. Nick's belief is clearly rational, and (in some good sense at least) what Nick believes is epistemically the *right* think to believe. Conversely, excellent but misleading evidence can make it appear that one epistemically ought *not* to believe a proposition which is in fact true. For example, in the above case, Nick should not believe that Gale's socks are green, even though they are. It would be irrational – an epistemic mistake – for Nick to believe that they are green, given that Gale has told Nick they are blue, and Nick has every reason to trust her and no other relevant evidence.

A set of general principles concerning epistemic good standing which attempts to respect these intuitions holds that claims of the form:

G1: S's belief that p is in epistemic good standing (or would be if S had such a belief)

are made true by the same states of affairs as claims of the form:

P1: p is probably true given S's state of information.

And, similarly, we can propose that claims of the form:

B1: S's belief that p lacks epistemic good standing (or would if S had such a belief)

are made true by same facts as claims of the form:

P2: It is not the case that p is probably true given S's state of information.

(As with truth, the naturalist who favours this sort of principle faces the task of naturalizing the required notion or notions of probability.)

Note that the 'probability' talk in these principles may be interpreted either as talk of objective chance or as talk of S's credence. The former will result in a less internalist notion, the latter in a more internalist notion. The talk of 'information' might also be interpreted in various ways, which differ concerning the extent to which one can tell from the inside what 'information' one has. The easier it is to tell from the inside what information one has, the more internalist this kind of epistemic normativity will be.

Even this second, more complex, set of principles might be thought to be too simplistic. It might be thought to run into difficulties, for example, in cases where one's state of information renders a belief that p highly probable but one believes p for reasons which are completely divorced from this fact. This suggests revising the principles to include a clause about the belief's *sensitivity* to the probability of its truth given one's evidence. Again, more and less internalist notions of sensitivity will be available.

But important though it is, this is not the place to continue the process of revising and problematizing particular proposals. For the points to be made below are entirely independent of which natural facts are to be associated with which readings of which

epistemically normative claims.

So far our discussion has focussed on the kinds of natural facts that one could associate with claims concerning epistemic good standing for beliefs. This doesn't settle the question of which facts should be associated with non-negated epistemic 'ought' claims. (Negated 'ought' claims could be associated with the same facts as claims of the form of B1.) It does not obviously follow from the fact that a belief would be in epistemic good standing if you had it that you epistemically ought to have it. In fact, any claim to that effect seems ridiculously strong.

There is the option of saying that it is epistemically permissible to have any belief which would be in epistemically good standing if you had it, but *no* belief is such that one ought to have it. But, plausibly, there are some situations where one epistemically ought to believe certain propositions. Perhaps the clearest cases are those where the evidence for a proposition is staring you in the face, and you are aware of this fact, and of the fact that there is no countervailing evidence. How will the naturalist account for this?

One option that sits well with the second set of principles discussed in this section is to say that claims of the form:

E1: S epistemically ought to believe p

are made true by the same states of affairs as claims of the form:

P3: p is probably true, in an obvious and salient way, given S's state of information.

Of course obviousness and salience are imprecise notions and await further explication, but the gist of the idea is enough for current purposes. More (or less) sophisticated proposals will be available for those who have quibbles with this one.

For now, it is enough if these examples have given the reader an idea of the relevant *kind* of naturalizing strategy to get a discussion underway of its merits relative to other strategies.

III Naturalizing Epistemology The Old-Fashioned Way

Attempts to naturalize epistemology are most often associated with Quine, who claims that epistemology is, or should be, merely ‘a chapter of psychology and hence of natural science’ (1969, p. 82). That is, he thinks epistemology is, or should be, a description of how we in fact think and reason.

However, in an attempt to reconcile this position with the apparent normative force of certain epistemic claims, Quine (at least sometimes) suggests that epistemic normativity is a species of instrumental normativity – a matter of what would be appropriate means to certain of our ends. He claims that normative epistemology is ‘the technology of truth-seeking ... a matter of efficacy for an ulterior end, truth’ (1986, p. 665). For this reason believes that ‘the normative here, as elsewhere in engineering, becomes descriptive when the terminal parameter is expressed’ (1986, p. 664-5). The envisaged ‘terminal parameter’ is the goal or end of finding out the truth, and the descriptive facts with which Quine thinks epistemically normative claims are to be associated (in some way or other) are facts about what is conducive to that end. For instance, the claim that one ought not to believe *p* is to be somehow associated with the fact that believing *p* is not conducive to one of one’s ends - namely, that of finding out the truth. Laudan (1987), and various others, take similar views. (See Kelly 2003 for an interesting and comprehensive discussion.)

Despite what Quine's writings suggest, it would be very surprising if things were in fact this simple. Again, it seems that misleading evidence cases can make belief in p fully epistemically correct (on at least one good sense of 'correct') even when p is false, and liable to give rise to belief only in further falsehoods, so that believing p is not conducive to the end of finding out the truth. Similarly, believing p truly but irrationally is conducive to finding out the truth about p but does not therefore seem to be epistemically permissible (at least on one good sense of 'epistemically permissible').

Be that as it may, a version of the Quinean view is compatible with the strategy of this paper. One could hold that epistemically normative claims pick out and are made true by facts about what is a means to the end of finding out the truth, even though they do not mean exactly the same as (have the same sense as) claims which are explicitly about such means-end relations. Other versions of the Quinean view, not compatible with the strategy of this paper, will hold that epistemically normative talk *does* mean the same as talk about means to ends, and still others will hold that all epistemically normative talk is literally false and should be *replaced* with talk about means to ends. It is not a straightforward matter to say which version of the view Quine and his various followers intend to defend.

A challenge to (all forms of) the Quinean approach is developed in Hookway 2000. Hookway argues that, once a particular 'terminal parameter' – a goal or aim of believing – has been identified, we still have a question as to whether the goal in question is one that we epistemically ought to pursue (see §9). Presumably Hookway's thought is that if nothing is said on this point, then it is unclear what bearing the fact that believing p advances or does not advance one's goals has on the question of whether believing p is

something one epistemically ought to do. Believing *p* is something one epistemically ought to do if and only if it advances the goals one epistemically *ought* to have: it is irrelevant what goals one actually has.

I'm not sure how effective this is as an objection to the Quinean view. One available response is that epistemic norms are norms concerning *belief*, and that other states, such as having an aim or goal, are not the sorts of things to which epistemic norms apply any more than they apply to desires or to actions like drinking tea, except perhaps by extension. (This sort of point is also made in Field 2000, p. 142.) At the very least, Hookway needs to do more to motivate the thought that it is reasonable to enquire as to the epistemic correctness of one's goals. A second response is that it seems question-begging against the instrumental view to assume that there is more to a belief's being epistemically correct than the fact that it advances certain of our goals.

However, regardless of its efficacy against Quine, it is worth noting that Hookway's objection has no mileage as an objection to most versions of the naturalizing strategy defended in this paper, including the two versions discussed in §II above. The defender of such a strategy need make no appeal to goals or ends in saying what the epistemically normative facts are like.

Note, though, that an appeal to goals or ends *might* enter the picture in a different way, which would not be susceptible to Hookway's challenge. For example, suppose it is characteristic of believing that in having a belief one aims to have a belief that is true (and/or that one aims to have a belief that is probably true given one's state of information). This could be appealed to in explaining the plausibility of proposals like those in §II. That is, it might help to explain why it is facts about (probable) truth, rather

than any other kind of facts, that are to be identified as the facts which make true epistemically normative claims.

Even on this view, however, the epistemically normative facts would be identified as facts about what is true (or probably true given one's state of information). Given this identification, it does not matter, for the purposes of determining whether or not S ought to believe p, whether or not S epistemically ought to have whichever goal it is that helps us explain why a particular set of natural facts is associated with epistemically normative claims, even assuming that it makes sense to assess a goal for epistemic correctness.

Nonetheless, one might wonder whether there isn't scope for an objection to any version of the strategy of this paper, similar in spirit to Hookway's worry, which runs along the following lines:

“Granted that p is true (or probably true given your state of information, or whatever), there is still a question whether your belief that p is epistemically in good standing, because there is still a question whether you epistemically ought to have beliefs which are true (or probably true given your state of information, or whatever)”.

This challenge is analogous to the Open Question Argument in ethics and similar challenges elsewhere. (See, for example, Smart's classic 1959, where he addresses comparable reactions to materialism about sensations.) By analogy with the response given above to the ethical Open Question Argument, one can respond here by saying firstly that, if all that is meant is that it is not *conceptually* and/or *obviously* true^v that a belief is in good standing iff it is true (or probably true given one's state of information, or whatever), then this is no objection to the thesis that the two propositions pick out and

are made true by the same fact. But if there is meant to be more to the objection than this, it is hard to see what it is supposed to be. To claim that the question of whether a belief is in good standing remains genuinely open once it is settled that it is true (or probably true given one's state of information, or whatever) is simply to beg the question against the view that the same fact makes true these two claims.

The strategy can also avoid a familiar problem for Quinean instrumental views (described clearly in Kelly 2003), which is that some people might simply lack the aim described in the 'terminal parameter', and hence fail to be subject to epistemic norms by the lights of the instrumental approach. For example, if S does not have the goal of finding out the truth then it is not the case that believing an obviously false proposition p is unconvincing to one of S's ends, namely, that of finding out the truth. Hence, according to (this version of) the instrumental view, it is not the case that S's belief that p is epistemically incorrect. This would be worrying, because clearly someone *can* be utterly epistemically irrational despite (perhaps because of) her failing to have the goal of finding out the truth. (And clearly, we can motivate a similar objection whatever other goal we might take to be the terminal parameter).

This kind of objection has given rise to sophisticated versions of the instrumental approach, due to Kornblith (1993) and Stich (1990), according to which everyone who has goals at all is bound to have the right ones for him to be subject to epistemic norms. But we can avoid the need for such subtleties. According to (most versions of) the strategy defended here, no-one need have any particular goals, or any goals at all, in order to be evaluable in epistemically normative terms. For instance, no-one need have the

goal of believing what is probably true given her state of information in order to be evaluable in terms of epistemic norms, if the low probability of p given S's state of information is enough to make S's belief in p epistemically incorrect. This point may be made clearer by analogy with the moral case: if the fact that X is morally wrong is simply the fact that X fails to maximize utility, no-one needs to be *aiming* to maximize utility in order for an action of theirs to be morally wrong.

IV Other Naturalizing Strategies

Goldman (1986) has made a proposal according to which certain epistemically normative claims (or what he takes to be epistemically normative claims, namely claims about justification – see p. 20) are made true by the same facts as certain natural-sounding claims. Yet his strategy is importantly different from that of this paper. He associates a class of natural-sounding claims that have not been mentioned here (namely, facts about reliability) with a class of (what he takes to be) normative claims that have not been central to our discussion (namely, claims about justification). But this difference is not important; as has been stressed, the strategy under consideration is compatible with various specific proposals about which natural facts to associate with which normative facts. What matters is the kind of association that Goldman proposes.

Goldman takes the relevant natural-sounding claims to *mean the same* as the corresponding normative claims. On p. 108, for instance, he says that 'the meaning of the term 'justified' (in its epistemic use) is fixed by certain things we *presume* about the world ... Specifically, beliefs are deemed to be justified when (roughly) they are caused by processes that are reliable in the world as it is assumed to be' (emphasis added). His

principle P1 on p. 59 is discussed as a candidate 'semantic truth about the language of justified belief' (and presumably the same goes for the improved principle P3 on p. 63 which Goldman endorses). He stresses at pp. 38-9 that his method for addressing scepticism is through 'conceptual analysis' of terms like 'justified', and it is clear that he takes this to mean analysis of the *meanings* of such terms, since he finds it relevant to defend the notion of meaning from Quinean attack. He relatedly thinks that our *concept* of justified belief is a concept as of belief caused by reliable processes; he says on p. 107, for instance, that 'our concept of justification is constructed against a backdrop of ... a set of normal worlds' (emphasis added). It is important to note that the naturalizing strategy of this paper is committed to no such claims about meaning, concepts or conceptual analysis. Many potential objections to Goldman – that this is not what epistemically normative terms mean, that conceptual analysis cannot deliver substantive results of this kind, etc. – are thus sidestepped.

Stich (1990) and Kornblith (1993) have objected to Goldman that even if his view is correct, we're not interested in what *we mean* by our epistemically normative terms, but in *what we epistemically ought to think* (see e.g. Kornblith 1993, p. 361) – a move that calls to mind a standard objection to Strawson's views on induction (Strawson 1952, pp. 248-63).

It is not clear that this is a fair move even against Goldman. (Its analogue also seems an unfair response to Strawson, for similar reasons.) Suppose, for instance, that it is established that what we mean by 'S ought to believe that p' is the same as what we mean by 'p is true'. Then that provides the resources to enable one to say something about what we epistemically ought to believe: namely, that we epistemically ought to

believe whatever is true. Anyone who doubts this conclusion must take issue with the synonymy claim; nothing short of that will enable him to resist the conclusion. The real problem with adopting a strategy like Goldman's is that it is problematic – and unnecessary for the project of naturalizing epistemic normativity – to suppose that epistemically normative claims have the same meanings as non-normative claims. In any case, even if there is anything to the worry Stich and Kornblith raise for Goldman, it will not affect the view defended in this paper, since the latter is not a view about what epistemically normative terms mean.

Kornblith also objects to Goldman that ‘semantic consideration alone ... cannot explain the normative force of epistemic terms’ (p. 362). There are two points to note about this type of objection. Firstly, it is not clear that the biconditionals should have to ‘explain the normative force of epistemic terms’ even if they are conceptual truths as Goldman suggests. The question of what – if anything – *explains* the normative force of ‘A’s belief that p is justified’ seems an entirely different question to that of whether this sentence means the same as ‘A’s belief that p is reliable’.

It remains to be seen, however, whether anything akin to conceptual analysis can deliver the truth of the biconditionals. The second point to note is that the relation envisaged in this paper between certain normative claims and their corresponding natural-sounding claims is consistent with the relevant biconditionals failing to be delivered by semantic consideration alone, or conceptual truths, or analytic truths, or a priori knowable truths or anything in that vicinity. Biconditionals like ‘G1 if and only if P1’ may, on the strategy defended here, be a posteriori, even if they are taken to be necessary. (They may be analogous to ‘Something is a sample of water if and only if it is

a sample of H₂O’.)

Another option for those who prefer the strategy of this paper to Goldman’s is to say that the relevant biconditionals *are* a priori, but this is not because their left-hand sides mean the same as their right-hand sides. This is not an extraordinary feature: many other biconditionals share it. For instance: ‘7+5=12 if and only if all bachelors are unmarried’ and ‘Carrie is self-identical if and only if Carrie is identical to something’. One might wonder whether one could tell a priori that the two sides of a biconditional *picked out the same fact* (as opposed to merely that they are true) unless this was a matter of seeing that the two sides had the same sense. It seems that a priori reflection on the meanings of two sentences which differ in meaning could lead one to the conclusion that both sentences pick out the same fact. For instance, it could be discoverable a priori that ‘Julius invented the zip’ picks out the same fact as ‘The actual inventor of the zip invented the zip’, without it following that ‘Julius invented the zip’ means the same as ‘The actual inventor of the zip invented the zip’. These two sentences differ in sense because the first employs a name whose reference is fixed by a definite description, while the second employs the description itself. But if we have a priori access to this reference-fixing mechanism then we may well (assuming a certain conception of facthood) be able to know a priori that the same fact is picked out in each case.

A few other naturalizing strategies which are – or at least could be – different from the one defended in this paper should be noted before concluding this section. Firstly, one could maintain that epistemically normative claims could only be made true by non-natural facts, but that all epistemically normative claims are false. A flat-out error theory

is possible here, as is a more mollifying fictionalist account whereby we have always intended our normative claims to be understood as prefixed by an ‘according-to-the-fiction’ operator. Secondly, one might say that epistemically normative claims are not factual but merely evaluative or expressive of certain non-cognitive states. Blackburn (1996), Field (2000 and 2005), Gibbard (2003), and Ridge (forthcoming) have defended versions of this sort of view. These two options, however, feel counter-intuitive enough to be counsels of despair: to be considered only if there is no defensible option which better respects our pre-theoretic judgements and uses of epistemically normative vocabulary. They therefore shall not be discussed further here.

A very interesting third option (or rather, group of options) is to adopt some sort of functionalism about epistemic normativity.^{vi} One might say, for instance, that the property of *being in good standing epistemically* is a functional property, to be cashed out in terms of the role played by that property (or: our word for that property) in folk and/or scientific theories of (say) truth and/or belief. To gesture in the direction of what this might look like, consider that, just as some hope to understand ethically normative properties such as moral goodness by understanding the role they play in guiding and regulating action, so we might hope to understand epistemically normative properties such as justification and rationality by understanding the role they play in guiding and regulating belief.^{vii}

If some characterization of the relevant roles could be found, we could begin a search for realizers. We might, for instance, decide that the property of *being in good standing epistemically* is realized by the property of *being probably true given one's state of information*. We would then have option of identifying the normative property with

the realizer property – which would make for a view that could easily be developed in a way consistent the claims defended in this paper, by saying that normative vocabulary has the sense of role-property vocabulary but refers to the realizers – or arguing instead that the normative property is a role property distinct from the realizer property. This latter option could still be recognizably naturalistic – and, indeed, could also be developed into a version of the view defended in this paper – provided that role properties are accepted as non-spooky (and that an appropriately naturalistic realizer can be found).

An argument against the identification of the normative property with the realizer property might be made by considering an epistemic analogue of Moral Twin Earth (for which see Horgan and Timmons 1991). On Epistemic Twin Earth, the belief-guiding (or whatever) role is played by a different natural realizer property than on Earth. The argument would rely on our having the intuition that the Epistemic Twin Earthers' epistemically normative terms pick out the same properties that ours do, suggesting that our terms do not pick out the realizer properties but rather the role properties.

There may be some resistance to the idea that the role-property view is consistent with the proposal that epistemically normative claims do not mean the same as the corresponding natural claims – in this case, claims about role properties.^{viii} Presumably the thought would be that the appropriate role property can be arrived at merely by reflecting on what our epistemically normative words mean (or on what our epistemically normative concepts are like, or something similar). But it is far from clear why we need accept this (in either the epistemic or the ethical case). Even if we think that a way to discover the appropriate role property is by 'Canberra planning' – isolating a set of

'platitudes' about our epistemically normative properties, and using them to formulate an appropriate open sentence which our realizer(s) will need to satisfy well enough and better than any rivals – it is far from clear that the relevant platitudes need be available through mere reflection on the meanings of our words or on our concepts. They might, for instance, be claims of a widely-accepted, but a posteriori, folk theory about what we should believe under what circumstances.

It is far from clear that any form of functionalism about epistemic normativity is appropriate, mainly because it is so hard to see what kind of roles might convincingly be said to characterize the relevant normative facts and/or properties. But there is undoubtedly plenty of scope here for development of a cluster of interesting positions, some of which could be serious rivals to the fact-identity proposal of this paper, some of which will be versions of it.

V Questions That Remain To Be Answered?

Resistance to the current proposal will no doubt be felt by those who consider that it leaves the following question unanswered: *Whence the normative force of epistemically normative facts, if they're identical with natural facts of the kind you have discussed?*

This question has one reading on which it asks: *What makes epistemically normative facts normative if you're right that they are identical with natural facts of the kind you have discussed?* The answer to that question is supposed to be obvious: their being identical with the right natural facts is enough to make them normative. The view under consideration is that those natural facts *are* normative facts. Asking what *makes* them normative is like asking what *makes* facts about pineapples facts about pineapples. That's

just what they are.

However, someone who asked the above question could be intending to raise an objection akin to the Open Question Argument, i.e. claiming that the following is still a good question to ask at this stage in the debate: *Granted that S's belief in p is true (or probably true given S's state of information, or whatever), where does the normative fact that S's belief is in good standing come from?* The answer to this question is meant to be that the normative fact comes from – or rather, *is* – the fact that S's belief is true (or probably true given S's state of information, or whatever). To say that the question of where it comes from is still open must be either to beg the question against the proposed identification of the normative fact with the natural fact, or to say that the normative question is still not *obviously* settled, or that it is *open for all that conceptual analysis can tell us*, both of which are consistent with the position advocated here and do not present any kind of problem for it.

One other thing those who are tempted to ask the above question might be getting at is the question of *why one should care* about epistemic norms on the view under discussion. Kornblith (1993, p. 363) says that 'an account of the source of epistemic norms must explain why it is that I should care about such things'. But this is not a reasonable demand. Consider the analogy with the ethical case: why must one have provided resources to explain why we should care about ethical norms in order to count as having given an account of their metaphysical source? Why couldn't there be a true and informative metaphysics of normativity which has no power to explain why we should care about the norms it deals with?

The demand looks even more suspicious when we think carefully about what is

being asked of us. The demand for an explanation of why one should care about a set of norms is a demand for an explanation of a certain normative fact (the fact that one should care about a certain set of norms). In each case, there will be a reading of the demand on which the normativity involved in saying one should care is of the very same kind as that of the set of norms under discussion. For instance, we can ask why it is morally required that we care about moral norms. One kind of answer to this sort of question, in the ethical case, is an obvious (and not very informative) one: because it is morally wrong not to care about moral norms. Other, more informative, kinds of answer may be available to those who possess suitable theories: e.g. because not caring about moral norms fails to maximise utility.

On the other hand, the demand could have a reading whereon the normativity involved in saying one should care is of a different kind from that of the set of norms under discussion. For instance, we can ask why is it epistemically required that we care about moral norms. In answer to this sort of question, it is fair to reply that it is far from obvious that the explanandum is true, and that the burden of proof is with the questioner to establish this before requiring anything of us.

In the epistemic case, it is not clear that the first type of demand is fair. Caring about epistemic norms is not a kind of belief, so it is controversial to presume that epistemic evaluations are even applicable to this activity. And as in the ethical case, the burden of proof is squarely with the questioner to establish that the second kind of demand is legitimate, i.e. to establish the truth of any explanandum of the form *It is X-ly required that we care about epistemic norms*, where 'X' is replaced by something other than 'epistemic'. The burden of proof is with the questioner to establish, for instance,

that it is morally required that we care about epistemic norms, or practically required, or whatever. (And even if something of this kind could be established, it is pretty clearly not what was intended in Kornblith's question.)

References

- Blackburn, S. 1996. 'Securing The Nots', in W. Sinnott-Armstrong and M. Timmons (ed.s), *Moral Knowledge? New Readings In Moral Epistemology*, Oxford University Press, pp. 82-100.
- Boyd, R. 1988. 'How To Be a Moral Realist', in G. Sayre-McCord (ed.), *Essays on Moral Realism*, Ithaca, New York: Cornell University Press, pp. 181-228.
- Brink, D. 1989. *Moral Realism and the Foundations of Ethics*. Cambridge University Press.
- Field, H. 2000. 'Apriority as an Evaluative Notion', in P. Boghossian and C. Peacocke (ed.s), *New Essays on the A Priori*, Oxford University Press, pp. 117-49.
- 2005. 'Recent Debates about the A Priori', in T. Gendler and J. Hawthorne (ed.s), *Oxford Studies in Epistemology*, Oxford University Press.
- Gibbard, A. 2003. *Thinking How To Live*. Cambridge, MA: Harvard University Press.
- Goldman, A. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Horgan, T. and Timmons, M. 1991. 'New Wave Moral Realism Meets Moral Twin Earth', in *Journal of Philosophical Research* 16, pp. 447-65.
- Hookway, C. 2000. 'Naturalism and Rationality', in L. Decock and L. Horsten (ed.s), *Quine: Naturalized Epistemology, Perceptual Knowledge and Ontology*, *Poznan Studies in the Philosophy of the Sciences and the Humanities* 70, Amsterdam: Rodopi, pp. 35-53.
- Kelly, T. 2003. 'Epistemic Rationality as Instrumental Rationality: A Critique', in

- Philosophy and Phenomenological Research 66, pp. 612-40.
- Kim, J. 1988. 'What is Naturalized Epistemology?', in *Philosophical Perspectives* 2, pp. 381-405.
- Kornblith, H. 1993. 'Epistemic Normativity', in *Synthese* 94, pp. 357-76.
- Laudan, L. 1987. 'Progress or Rationality: The Prospects for Normative Naturalism', in *American Philosophical Quarterly* 24, pp. 19-31.
- Moore, G.E. 1903. *Principia Ethica*. Cambridge University Press.
- Quine, W.V.O. 1969. 'Epistemology Naturalized', in *Ontological Relativity and Other Essays*, New York: Columbia University Press, pp. 69-90.
- 1986. 'Reply to Morton White', in L. Hahn and P. Schlipp (ed.s), *The Philosophy of W. V. O. Quine*, La Salle, Illinois: Open Court, pp. 663-5.
- Railton, P. 1986. 'Moral Realism', in *Philosophical Review* 95, pp. 163-207.
- Ridge, M. Forthcoming. 'Epistemology for Ecumenical Expressivists', to appear in *Aristotelian Society Supplementary Volume* 81 (2007).
- Smart, J.J.C. 1959. 'Sensations and Brain Processes', in *Philosophical Review* 68, pp. 141-56.
- Stich, S. 1990. *The Fragmentation of Reason*. Cambridge, MA: MIT Press.
- Strawson, P. 1952. *Introduction to Logical Theory*. London: Methuen.
- Sturgeon, N. 1985. 'Moral Explanations', in D. Copp and D. Zimmerman (ed.s), *Morality, Reason and Truth*, Totowa, New Jersey: Rowman and Allenfeld, pp. 49-78.

Notes

* Thanks to Justin Tiehen, Jussi Suikkanen, Daniel Nolan, Neil Tennant and two anonymous referees for helpful comments and discussions.

ⁱ So Kim is right to say that there is room for a kind of epistemological naturalism which 'does not include ... the claim that epistemic terms are definitionally equivalent to naturalistic terms' (p. 398). However, Kim's alternative strategy (defending naturalism by appealing to supervenience) is quite different from the one to be defended here.

ⁱⁱ One way of defusing the Open Question Argument is to argue that even if there *is* a conceptual entailment between M1 and O1, it is extremely unobvious. The suggestion under discussion here is rather that there is no, even unobvious, conceptual entailment.

ⁱⁱⁱ It might be felt, for instance, that I1 is true on its externalist reading but false on its internalist reading. The latter intuition might derive from the thought that it is sometimes epistemically acceptable, in the internalist sense, to believe inconsistent propositions which are – with good reason – taken to be consistent.

^{iv} Such divergence of usage could be taken to parallel two apparently different uses of morally normative vocabulary, on the first of which one ought to do whatever *in fact* maximizes utility, while on the second one ought to do whatever maximizes utility *as far as one can tell*.

^v One way of responding to this analogue of the Open Question Argument is to say that there is a relation of conceptual entailment between certain naturalistic claims and the corresponding normative claims (cf. note i) but it is extremely unobvious. This strategy is not pursued here, mainly because it is appealing to think, in the absence of countervailing evidence, that reflective, competent users of a concept are normally pretty good at detecting conceptual entailments of such low complexity.

^{vi} This section owes much to helpful comments from Justin Tiehen, Jussi Suikkanen and Daniel Nolan.

^{vii} Interestingly, analogues of ethical internalism and externalism are available here. Just as an ethical internalist thinks that judging something to be wrong is intrinsically and/or necessarily connected with being motivated not to do it, so her epistemic analogue would think that judging a belief to be irrational

is intrinsically and/or necessarily connected with feeling some pressure not to have that belief.

^{viii} After all, the Horgan and Timmons thought experiment is employed to caution against ethical analogues of this sort of view. However, it is important to bear in mind that they have specific versions of synthetic moral realism in their sights, namely versions on which the referents of moral terms are more like realizer properties than role properties.