

Scientific Naturalism and the Value of Knowledge

Philosophical naturalism is, arguably, the dominant philosophical tradition in contemporary western philosophy. Naturalistic theories abound in nearly every area of philosophical investigation, and epistemology is no exception. Just what counts as a naturalistic theory in epistemology is not completely obvious, but the call for and interest in such is unquestionable.

When we begin to ask which theories count as naturalistic ones and which do not, things get more complicated, for it is far from obvious why any extant epistemological theory is anti-naturalistic. In the first section below, I will argue for a particular understanding of naturalism, situating the contemporary interest in it in a foundational attitude of respect for science. This motivation for the view constrains the kind of epistemological theory one can adopt, and my goal is to show how these constraints push inexorably toward a kind of attitudinalism or non-cognitivism in epistemology, an attitudinalism modeled on non-cognitivist approaches in ethics. The force doing the push is the need to account for the value of knowledge, which I will provisionally assume. My ultimate goal is to show how difficult it is for naturalistic views to account for the value of knowledge.

I. Naturalism and Respect for Science

Richard Feldman has argued recently that everyone, or nearly everyone, in epistemology can legitimately claim to be offering a version of naturalized epistemology.^[i] The difficulty of classifying any view as non-naturalistic is, at bottom, the cost of allowing non-reductive versions of naturalism. Consider, for example, what Scott Sturgeon says about mountains:

I'm a realist about [mountains]. But I'm no reductionist. I do not think a tidy condition can be stated, in non-mountain terms, which coincides with being a mountain. So I deny mountainhood is identical with such a condition. Does that make for trouble? Does realism about mountains interact with non-reductionism to force spooky mountain metaphysics?

Surely not. Mountains are nothing over and above non-mountains. I do *not* say that, though, because I can reduce mountains to non-mountains. I say it because the best picture of the world says it; and I adopt that picture. I say it because the perspective which strikes the best balance of

simplicity, strength and the need to resolve Sellar's Problem [concerning the relationship between the manifest and scientific images] *contains* the claim that there are mountains which are irreducible yet not fundamental. The best picture of the world entails non-reductive naturalism about mountains. It says mountains are non-reductively made of non-mountains. So that is what I accept.^{ii[ii]}

The difficulty for finding any non-naturalist theory of justification follows immediately:

And so it could be with justification. I'm a realist about it too. But I take seriously the possibility that justification does not align with a tidy descriptive condition. And if it doesn't not, justification isn't identical to any such condition. Would that make for trouble? Would realism about justification plus non-reductionism *force* spooky normative metaphysics? I don't see why. Perhaps the best picture of the world contains the claim that beliefs are justified in a metaphysically derivative but irreducible way. Perhaps the perspective which strikes the best balance of simplicity, strength and the need to resolve Sellars' Problem *endorses* non-reductive naturalism about reason. In the event, we should too. We should accept reason as metaphysically derivative yet irreducible. We should accept non-reductive naturalism about reason.^{iii[iii]}

The most popular way to articulate Sturgeon's naturalistic vision for reason and justification is in the language of supervenience, allowing us to say that normativity depends on natural facts without being reducible to them. Once we recognize such a naturalistic possibility, however, it is exceedingly hard to find any epistemological theory that cannot legitimately claim to be a naturalistic one. Even theories formulated exclusively in normative terms can endorse non-reductive naturalism by appealing to supervenience relations.

Feldman argues an even stronger point. Since it is far from clear what counts as a naturalistic concept at all, it is far from clear why an appeal to the language of normativity itself is somehow incompatible with naturalism. Chisholm, for example, thinks that there are synthetic a priori truths linking experience with the rationality of certain beliefs,^{iv[iv]} and Feldman argues that even such an appeal need not be incompatible with naturalism:

Chisholm apparently thought that, in addition to deductive and probabilistic connections, there was another species of connection between propositions (or between experiences and propositions). His view was that these relations are part of the real, or natural, world. Some may deny that there are any such relations. This seems, once again, to be a dispute about what there is,

not a dispute about whether there is something beyond what is natural. In other words, if Chisholm is right, it is quite unclear why terms such as "supports" and other epistemic terms do not belong on the list of naturalistically acceptable terms in the first place. If they do, then even Chisholm can plausibly maintain that epistemic support facts are natural facts. If so, then almost all epistemologists are substantive naturalists.^{v[v]}

I think Sturgeon and Feldman are onto something very important here, and it is that no one has a very good idea of what substantive naturalism is committed to, a point argued persuasively by Bas van Fraassen.^{vi[v]} He argues that if we specify substantive conditions for being a naturalist, we need only wait around for a few decades and the developments of science will show that our view is mistaken. What happens as a result is not a rejection of substantive naturalism, but rather a change in the account of what the view involves. This practice van Fraassen labels "false consciousness" since it shows that there is some underlying viewpoint or attitude responsible for the practice of revising the account. Whether or not we agree with van Fraassen on this point, it is important to distinguish between substantive versions of naturalism and methodological ones, where a methodological version of naturalism arises from respect for science and its methods of investigation. Since science is an investigation of the natural world, any viewpoint that gives honor to the methods of science in the search for truth will incline toward naturalism. Full endorsement of naturalism results when one maintains that no truth can be beyond the epistemic reach of the methods of science, that everything that exists can in principle be investigated by these methods. Some will also add the claim that the only respectable methods for investigating what exists are the methods of science, but naturalism does not require this stronger claim. It is enough that the long arm of science has the entirety of truth within its reach.^{vii[vii]}

Given this construal of naturalism, the attempt to provide naturalistic theories in epistemology is more constrained than we were able to derive above. If one's goal is to provide a naturalistic epistemology and we have only methodological naturalism as our guide, then one can be sure of the naturalistic status of one's theory by employing only those concepts that are already needed in a scientific description of the world. So, for example, causal theories of knowledge are paradigm instances of naturalistic theories in epistemology, since causality has a presumptive claim to being needed to give an adequate account of the natural world.^{viii[viii]} So too do most versions of reliabilism, according to which beliefs are assessed as candidates for knowledge in terms of nature of the processes or methods that underlie the belief, specifically, whether those processes or methods generally lead to true

beliefs. Here the concepts necessary to express the view involve concepts surely at home in the sciences: concepts such as processes and probabilities.

Most versions of virtue epistemology have the same grounds for claiming to be naturalistic theories, inasmuch as they identify the virtues in terms of stable characteristics of a cognizer that yield true beliefs most of the time. Similarly, Plantinga's proper function theory of warrant appears at first glance to have equal claim to being a naturalistic theory, in virtue of the necessity of employing the concept of a proper function in the biological sciences.

This last example reveals a caveat I ought to note about the entitlement to substantively naturalistic credentials on the basis of employing conceptual apparatus central to contemporary science. Even though biology is rife with appeals to proper function, it may be that the concept is implicitly non-natural nonetheless, and Plantinga in particular argues that it is. He argues, that is, that the concept of proper function implies that of a design plan and of a designer (beyond that of Mother Nature).^{ix[ix]} Thus, even though his theory is formulated in terms of concepts central to contemporary science, that fact is no guarantee that the theory will be substantively naturalistic.

Even so, constructing a theory using scientifically respectable concepts is the best path to follow for methodological naturalists, since following this path warrants the naturalistic credentials of one's theory. This approach to what counts as naturalistic has the virtue of leaving a large number of theories lacking naturalistic credentials, theories such as coherentism, foundationalism, evidentialism, and some versions of contextualism. What unifies this variety is that each such theory can be expressed employing the concept of evidence (without explaining that concept itself in naturalistic terms): coherentists typically have a holistic conception of evidence, foundationalists typically require the existence of basic evidence, as do contextualists of a certain variety, though they allow the class of basic evidence to vary by circumstance (and evidentialists, most obviously of all, must employ the concept of evidence to express their view).

II. The Problem of Naturalistic Purity

Once naturalism is understood in this way, it is easy to see why naturalistic epistemologists gravitate toward various versions of reliabilism. Such approaches employ conceptual apparatus clearly at home with the

language of science, and they can do so in a way that requires no mysterious supervenience claims about the dependence of the normative on the non-normative.

Even so, endorsing such an approach is no guarantee of producing a naturalistically acceptable theory of knowledge. To see the problem, recall that theories which appeal to evidentialist language are naturalistically suspect. The problem is that naturalistic theories have grave difficulty avoiding such an appeal; that is, they have a difficult time maintaining what I will term “naturalistic purity.”

Some examples may help illustrate how a theory might fail the test of purity. According to evidentialism, justification is required for knowledge, and the nature of justification is to be understood in terms of that which is supported by one’s evidence. Of course, not every belief for which one has evidence is justified, for one can have evidence for a claim and yet have grounds for doubting it that are sufficient to warrant withholding. A pure version of evidentialism will attempt to explain this grounds-for-doubt qualifier in terms of the concept of evidence itself, and it is not hard to see how to do so. A ground for doubt, according to evidentialism, is a piece of information possessed by a cognizer which is such that in conjunction with any evidence a person has for a certain claim, fails to provide adequate evidence for that claim. So evidentialism seems to have little difficulty passing the purity test: it employs the same conceptual resources to explain the grounds-for-doubt clause as it does to clarify initial, or *prima facie*, justification.

Not so for reliabilism, however. One might initially think that a belief is candidate for knowledge, according to reliabilism, when it is produced or sustained by a reliable process or method. One difficulty with such a proposal is that one may have grounds for doubting the claim that are sufficient to undermine it as a potential candidate for knowledge. So candidacy for knowledge requires the absence of certain types of grounds for doubt, and a pure version of reliabilism will try to clarify this condition appealing to the language of reliability.

Alvin Goldman’s 1979 discussion of this issue is instructive.^{x[x]} He attempts initially to deal with such cases in terms of alternative, available reliable processes that would not have resulted in the belief in question. Goldman immediately recognizes that this approach won’t quite do as it stands. He points out that requiring an additional process to be used can’t be quite right—if another process is used, the first one won’t have been used at all. He also worries about the concept of availability, wondering exactly what it takes for a process to be available (was the scientific method, for example, available to Plato?).

In response to these problems, Goldman gives up the task of constructing a pure version of reliabilism. He says, "What I think we should have in mind here are such additional processes as calling previously acquired evidence to mind, assessing the implications of that evidence, etc."^{xix} I think Goldman is issuing a promissory note regarding some analysis he is not yet in a position to construct. I think, that is, that he wants to construct a pure version of reliabilism, one which does not borrow language from alien theories of knowledge. Since he doesn't see how to do so, he characterizes the processes in terms of the concept of evidence, leaving grounds for complaint from naturalists. For an impure version of reliabilism is not obviously acceptable to naturalists, and one doesn't get a pure version of reliabilism simply by appending the phrase "the process operative when" to the key explicatory concepts of an alien theory.

Plantinga's proper function theory provides another example of the same impurity. The key concept of his theory is that of proper function, and he expresses enjoyment at being able to provide such a naturalistically-acceptable theory of knowledge.^{xii} In response to the need to rule out possible grounds for doubt, however, Plantinga introduces the concept of a defeater and that of a defeater system, and makes no attempt to explain these concepts in terms of that of proper function.¹³ Moreover, the concept of a defeater is pretty clearly an evidential concept, no matter what kind of defeater is involved.¹⁴ So Plantinga provides only an impure version of a proper function theory of knowledge.

Plantinga, of course, holds no stake in a defense of naturalism, for he is convinced that naturalism is false and that the notion of proper function is not, at bottom, compatible with naturalism anyway. For those who thought they had found an ally of naturalism in an epistemology of proper function, however, the point is instructive that Plantinga's version, impure as it is, offers no aid and comfort to his enemies.

Virtue epistemologies of a reliabilist sort face the same problem. The need for a grounds for doubt clause is pressing here as well. Suppose, for example, that one's perceptual faculties count as intellectual virtues or excellences. It is still possible for one to employ such faculties when one's background information argues against using them. In such a case, the temptation Goldman succumbed to presents itself: add a ground for doubt clause formulated in evidentialist language, and thereby insure a failure of naturalistic purity.

Perhaps, though, one might try to avoid the problem as follows. One might think of the problem as a special version of the problem of generality,¹⁵ and the solution to it one of specifying a field of operation in which

the use of the faculty always implies positive epistemic status. Consider Sosa's proposal, for example. Sosa relativizes the notion of a virtue to conditions, fields, and environments. He says,

One has an intellectual virtue or faculty relative to an environment E if and only if one has an inner nature I in virtue of which one would mostly attain the truth and avoid error in a certain field of propositions F, when in certain conditions C. The distinction between E and C is not sharp or important and amounts to a distinction between relatively stable background conditions and relatively episodic conditions.¹⁶

When Sosa gives his final account of a virtue, he supplements the above rough account with a further condition that no broadening of C, F, and E are such that the person would be likely to hold a belief with respect to *p* and not be likely to be correct regarding *p*.¹⁷

It is worth noting in passing that this account is not quite what Sosa needs, for it succumbs to an evil demon problem. We can define evil demon counterparts for C, F, and E, and then specify a broadening of each in terms of the disjunction of each with their respective evil demon counterparts. Then for cases in which a person is inclined to form a belief and would still be so inclined in a demon world (perceptual beliefs, for example), then there is a broadening of C, F, and E in which the person is likely to form a belief but is not likely to be correct.

Let us put aside this difficulty, however, for we are not here concerned with the problems raised in epistemology by the possibility of demon worlds, but rather with the question of whether some version of virtue epistemology can pass the test for naturalistic purity. What is needed in Sosa's account for that to be so is for the reference to E, C, and F and the broadening of each to be sufficient to rule out cases where, e.g., a perceptual belief is likely to be correct, even though the believer has background information that makes forming perceptual beliefs in those circumstances suspect. So suppose a person is in such circumstances, i.e., suppose there are values for E, C, and F such that S believes that *p* in <E,C,F> and is likely to be correct in such circumstances. Suppose further, however, that S has background information casting doubt on the wisdom of forming any such belief in those circumstances.

Sosa's account implies that the belief nonetheless has positive epistemic status, and his response to such cases would appeal, I think, to his distinction between animal and reflective knowledge and his related distinction between aptness and justification.¹⁸ Justification is a matter of standard coherence relations on bodies of information, whereas aptness is a matter of having been produced by a virtue, allowing the distinction between

animal knowledge, which requires only aptness of belief, and reflective knowledge, which requires justification. We need not devote much attention to the details here, however, for the only way this approach can pass the test of naturalistic purity is to focus exclusively on the concept of animal knowledge. Once the coherence relations involved in Sosa's account of justification are brought into the picture, which includes logical, evidential, and explanatory connections within a body of information, it is clear that the theory has failed the test of naturalistic purity. Furthermore, even if we focus exclusively on the concept of animal knowledge, granting that there is a kind of knowledge that animals possess that need not involve the kind of justification Sosa delineates, it is far from clear that the same concept is applicable to human beings with sophisticated understandings of their epistemic situation, aware of grounds for being suspicious of the reliability of a type of belief-forming mechanism and nonetheless ignoring those grounds for doubt. That is, even if we wish to posit some truncated concept of knowledge appropriate for small children and animals, who lack the sophisticated perspective on themselves that is distinctive of reflective knowledge, we should not extend the application of that concept to those having such a reflective perspective and who choose to ignore its implications.

If we alter the account to accommodate this point, Sosa's theory has no resources to avoid an intrusion by evidentialist concepts into the basic concept of animal knowledge, for that account will need to include a rider to the effect that if one has a reflective perspective on one's situation, then one will need to lack information casting doubt on the wisdom of using one's belief-forming mechanisms in that situation.

Here John Greco's version of virtue epistemology is helpful and illuminating, for Greco attempts to provide a thorough-going virtue account of a problem akin to the one we are addressing. Greco grants the need for both objective and subjective justification in an account of knowledge, and offers a relatively straightforward virtue account of objective justification. What is unique in his view is that he also attempts to provide a virtue account of subjective justification, and if he were successful in this project, we would have a way of developing a virtue account that passes the test for naturalistic purity.

According to Greco, objective justification amounts to a belief being the result of dispositions that make a person reliable regarding that belief in the conditions in question, and subjective justification involves a belief being "the result of dispositions that S manifests when S is thinking conscientiously."¹⁹ Greco identifies conscientious thinking with the "default mode" of being motivated to get to the truth,²⁰ thus contrasting honest thinking with think that is motivated by non-alethic factors such as greed, prestige, comfort, and the like.

Second, Greco does not identify subjective justification with such proper epistemic motivation; in fact, his definition does not even require proper motivation. For his definition only requires the activity of the dispositions that are present when one is properly motivated, and those dispositions might be active both when one is properly motivated and when one is not. I think this may be too weak a connection between subjective justification and proper epistemic motivation. For example, if the same disposition can accompany both well-motivated and ill-motivated belief, it will be possible to display that disposition ill-motivatedly while rationally believing that an ill-motivated display on this occasion is highly unlikely to get one to the truth. Such a situation strikes me as paradigmatic for lack of subjective justification, rather than one in which subjective justification is present, as Greco's theory implies.

Moreover, Greco's appeal to subjective justification does not explain away the standard counterexamples to simple reliabilism. In Bonjour's clairvoyant case,²¹ Greco needs to charge the clairvoyant with not manifesting the dispositions that he manifests when trying for the truth. I don't see why the clairvoyant has to be guilty of this charge. The clairvoyant knows better than to trust clairvoyance, but that doesn't imply that he is manifesting dispositions different from the ones manifested for him when aiming for truth. The clairvoyant's failure is that he doesn't take into account possessed defeating information, but one can fail to do this, and even fail to be disposed to do this, and nonetheless manifest the dispositions one normally does when proceeding honestly.

Most of us do not behave this way cognitively, but that is a contingent fact about us. As we improve cognitively, we learn to monitor for defeating information and we learn to withhold belief when we learn of the presence of defeating information. Even so, in the process of so improving, we often think honestly and display the dispositions that we ordinarily display when honestly trying for the truth without monitoring for defeating information and without withholding belief in the presence of known defeaters. Greco's response to the clairvoyance case requires that the only explanation for retaining a belief in the presence of known defeaters is that we are being moved by dispositions other than those operative when thinking honestly, but there are other options. Habits are often overly general, displaying themselves even where not especially useful or desirable. Transparently honest people sometimes hurt others' feelings by unthinkingly displaying such honesty, and belief formation can exemplify this same feature. The motivations present when one is honestly trying for the truth might be unthinkingly displayed when attention to the presence of known defeaters would have prevented the display of these motivations.

This point is not merely a trifling failure of detail, but rather a particular instantiation of a more general weakness of reliabilism. On the face of it, knowledge and justification are functions of the information or evidence we possess, but reliabilism wishes to talk in terms of belief-forming mechanisms and so must try to mimic evidential relations with these mechanisms (say, by individuating mechanisms or character traits in such a way that the reliability of these mechanisms coincides with the intuitive idea of information or evidence possessed). I have argued repeatedly that the prospects for successful mimicry are hopeless,²² and my point here is that the above difficulty is simply another example of the same problem. The concept of defeating information is, intuitively, one concerning the epistemic relationships between semantic contents, and the hope of any version of reliabilism is to mimic these relationships by appeal to the right kinds of mechanisms or character traits. Greco's attempt on this point is not entirely successful, I think, for there are no grounds for thinking that in order to display dispositions involved in trying for the truth, one must be disposed always to avoid belief when defeating information is present. So there is no reason to think that the clairvoyant case is explained away by an appeal to such dispositions.

The point of the discussion to this point has only been a kind of sabotage of the naturalistic program, arguing as I have that it is not obvious what counts as naturalized epistemology and it is far from clear that naturalists have taken seriously enough the requirement to propose a fully naturalistic theory. I want to attack the position more deeply, however, for there is a further problem that shares equal billing among the ignored problems for naturalized epistemology: the problem of accounting for the value of knowledge. I will argue that if naturalists wish to retain the assumption that knowledge claims have truth value, they will be pushed in the direction of virtue epistemology in order to account for the value of knowledge. Since we have already seen that it is not obvious how to be a virtue epistemologist and a naturalist at the same time, the temptation will arise to jettison the assumption. That is, when faced with the difficulty of accounting for the value of knowledge, one's commitment to naturalism may lead one to the extreme position of denying that knowledge claims have truth value, much as a similar naturalistic outlook has led to a similar, non-cognitivist approach in ethics. After goading the naturalist into this position, I will argue that such non-cognitivism in epistemology is untenable.

Naturalistic Epistemology and the Value of Knowledge

The problem of the value of knowledge is largely ignored and underappreciated in the history of epistemology. The problem is first introduced by Meno in Plato's dialogue by the same name, where Meno tells Socrates that knowledge is valuable because it is the appropriate guide to action. Socrates provides a counterexample by way of response: if you want to get to Larissa, a guide with true opinions about how to get there is every bit as good as a guide with knowledge. Meno at first balks at the counterexample, but then upon seeing its force, replies, "In that case, I wonder why knowledge should be so much more prized than right opinion, and indeed how there is any difference between them."²³

Meno's response is instructive. Meno's original conviction was that knowledge is valuable and on pragmatic grounds, but on seeing Socrates' counterexample, he wonders whether our common view that knowledge is valuable is mistaken. His confusion here leads to perplexity about the nature of knowledge—perhaps, he imagines, knowledge is nothing more than true opinion. In Meno's mind, there is an interplay between the question of the value of knowledge and the question of the nature of knowledge, so that failed assumptions about the value of knowledge generate doubts about one's assumptions regarding its nature.

I believe Meno is onto something. Though the question of the nature of knowledge is a prominent and important part of the history of epistemology, the enterprise itself presupposes the value of that about which we theorize. Aristotle reports that all people by nature desire to know, and in theorizing about the object of this desire, we should our presumption that it is a desire for something important. Thus, even though the history of epistemology has not devoted much attention to improving on Meno's failed pragmatic account of the value of knowledge, there is good reason to view the question of the value of knowledge as central to the discipline. We may put the point this way: an adequate epistemology is presumptively subject to twin desiderata, one concerning the nature of knowledge and the other concerning the value of knowledge. First, no epistemology that misconstrues the nature of knowledge can be adequate, and second, an epistemology that undercuts any attempt to explain the value of knowledge is defective.

I said above that an adequate epistemology is "presumptively" subject to these desiderata. This qualifier applies primarily to the requirement concerning the value of knowledge. I do not wish to claim that no epistemology can be adequate without containing an account of the value of knowledge. Nor do I claim that an epistemology is inadequate simply because it is incompatible with the value of knowledge. I do claim, however, that such incompatibility is a strong reason against the adequacy of such an epistemological theory. Such an

implication regarding the question of the value of knowledge would have to be accompanied by an explanation of why knowledge is not valuable.

As I have said, the history of epistemology contains little discussion of the question of the value of knowledge. Given the explicit recognition of the significance of the question in Plato, this failure is somewhat surprising, but I think an explanation for it is not hard to find. The concepts by which knowledge is distinguished from true opinion—concepts such as justification, certainty, infallibility, adequate evidence, accompaniment by an account or reason, having the right to be sure, etc.—are all obviously evaluative concepts, and positive ones at that. As such, it would seem a bit pedantic to spend much time arguing that they add something of value beyond that which true belief provides. So the issue of the value of knowledge remains at the level of a presupposition, in no need of much explicit discussion. It is secure in that position in virtue of the selection of approaches to the question of the nature of knowledge that cite evaluatively positive properties.

What is interesting to note in the present context is the special difficulty faced by naturalistic theories regarding the value of knowledge. The problem that such versions of naturalism face I will term “The Swamping Problem.” It arises at the most abstract level because of the following possibility. It is possible for property P to be valuable and for property Q to be valuable, and yet for the property of being both P&Q to have no more value than the value of one of its components. In such a case, the value of one of the components swamps the value of the other component.

The swamping problem would not be relevant to the question of the value of knowledge if that question were merely the question of whether knowledge has value. It follows that knowledge has value if it is composed of items some of which are valuable and not of which are not. The difficulty is that the question of the value of knowledge, as raised by Socrates, is not merely the question of whether knowledge has value. It also involves the issue of whether is more valuable than related cognitive states such as mere true belief. Once we raise the question of whether knowledge is more valuable than mere true belief, the possibility of encountering the swamping problem is raised since the value of true belief may swamp whatever value there might be in other features of knowledge.

Such a problem is most obvious when the additional item is likelihood of truth. For since true belief is valuable, it is clear that the value of likelihood of truth is parasitic on the value of truth itself. The property of being likely to be beautiful is a valuable feature of a painting, but only because being beautiful is itself a valuable property

of a painting. Moreover, if a painting is known to be beautiful, no additional value is judged to be present by noting that it is also likely to be beautiful. In a phrase, the value of the latter is swamped by the value of the former.

One more example to drive home this point. Suppose I want chocolate, and also want to expend little effort in finding it. I run a search on the web and find a list of places within walking distance of my campus that sell chocolate. I find another webpage that gives a list of places within walking distance that are likely to sell chocolate (the methodology for constructing this list is not given, but perhaps if a place sells any kind of candy it is likely to sell chocolate, and so it is a list of places that sell candy, though not presented as such). The first list is of more interest to me than the second, but that is not my point here. Instead, my concern is with a third webpage, one which is the intersection of the first two: a list of places that both sell chocolate and are likely to sell chocolate. This third list is of no more interest to me than the second, given that my goals are only to find chocolate within walking distance. Moreover, the third list is of no more interest to me than the first, even though I have an interest in places that sell chocolate over those that do not and I have an interest in places that are likely to sell chocolate over those that are not likely to sell chocolate. That is, even though both properties are valuable from my perspective, the value of one swamps the value of the other so that the combination of both properties does not yield any enhanced value.

What happens to the swamping problem if one eschews naturalism in epistemology? Put simply, the problem is not severe, for there are several options available for addressing it. First, one might hold that there is a special constituent of knowledge that is intrinsically valuable on its own, independently of any relationship it has to the truth.²⁴ Another option is to turn more subjective, and identify justification with doing the best by one's own lights in getting to the truth and avoiding error. Such an account of justification explicates its value in terms of its relationship to truth, but doing one's best can have value independently of increasing the likelihood of truth in any way at all. Hence, the value of doing one's best in trying for the truth is not swamped by the presence of truth.

None of this is meant to imply that every theory of knowledge that does not self-consciously aim at being a naturalistic theory of knowledge will be immune from the swamping problem. What I am arguing, instead, is that naturalistic theories of the reliabilist sort face a special burden, which leads to the question of whether there is a special variety of reliabilism that avoids this problem. Virtue epistemologists think there is, for they think that the analogy between virtuous belief and virtuous action is sufficient to undergird the claim that one deserves credit for virtuous belief just as one deserves credit for virtuous action.²⁵

Such a strategy must face the issue raised by the role autonomy plays in deserved credit for virtuous action. Actions are voluntary, whereas beliefs are not, and it is quite natural to think that this distinctive feature of action undergirds the practice of praising and giving credit for certain actions.

There are deep and important issues here that I am going to pass over in order to pursue the remainder of the discussion proposed on the topic of naturalism. So instead of pursuing the details, let me gloss quickly what I think a virtue epistemologist needs to say. The first point to insist on is that credit due for action does not require a libertarian conception of freedom to undergird the voluntariness central to any case of credit due for a good action. Instead, credit can be due when it is produced by the right kind of internal cause. In the case of credit due for action, it is important that the internal causes are such that voluntariness occurs. In the case of credit due for belief, it is also important that the belief is the result of internal causes, and even though the right kind of internal causes need not produce beliefs which count as voluntary ones, credit due nonetheless depends on the presence of these causes. In this way, one can deserve credit for certain types of belief just as one can deserve credit for certain types of actions. What is central to the analogy is the internal nature of the causes, not the concept of voluntariness.

There is much more to be said on this issue, but I will leave this discussion with the following provisional conclusion. The issue of the value of knowledge is an important one, and it is worth noting how restrictive a desideratum it is, once one embraces the project of naturalizing epistemology. If knowledge is truly valuable, the only direction I can see for the naturalist to go while still embracing cognitivism is to virtue epistemology.

The problem is that virtue epistemology is, for the naturalist, a house built on sand. When the storms due to questions about naturalistic purity arise, the building is likely to collapse. For no virtue approach appears capable of refraining from appeal to concepts that the naturalist cannot tolerate. In light of this fact, what is a naturalist to do? The temptation is, I think, to turn non-cognitivist about epistemology just as logical positivists turned non-cognitivist about ethics. When one's basic philosophical orientation cannot accommodate a certain domain of truth, one of the two has to go. Thus arises the temptation toward non-cognitivism in epistemology, and it is instructive and not surprising to find in recent discussion the first steps toward such a picture in epistemology. After all, if one of the key problems in epistemology is accounting for the value of knowledge, such theories seem well-suited to explaining that value immediately and directly in terms of the attitudes expressed in using such language. Moreover, the difficulties faced by the cognitivist in preserving naturalistic purity while attempting to account for the value of knowledge are so difficult that it might seem the path of least resistance to treat at least some of the

subject matter along the attitudinal or expressivist lines that are so popular in ethics and value theory more generally. So we can think of the discussion to this point as aimed at pushing the naturalist toward such an attitudinalist or expressivist position.²⁶ In the remainder of this essay, I want to explore briefly the prospects for this approach.

Attitudinalism and the Value of Knowledge

On the face of it, attitudinalism is perfectly adapted to handling the problems for cognitive versions of naturalism discussed above. The technical details of how to handle a grounds for doubt clause disappear, since questions of cognitive content disappear in favor of attitudes expressed. And the question of the value of knowledge seems to be addressed directly, since the attitude expressed in attributions of knowledge is a positive one. There is therefore reason for optimism that attitudinalism provides a solution to the problems facing the naturalist in epistemology.

Attitudinalism can be developed in differing degrees of scope. The widest scope would be an attitudinalism that covered the entirety of knowledge, claiming that the concept was entirely non-cognitive in nature. A weakness of this version of the view is that one the factive character of knowledge, that one cannot know and be mistaken. This point prompts an interest in attitudinalisms of lesser scope, ones that attitudinalize the normative dimensions of knowledge or those that attitudinalize the entirety of the difference between knowledge and true belief.

Concerns about naturalistic purity arise again if one attitudinalizes only part of the difference between knowledge and true belief. For example, suppose one adopts attitudinalism about justification, as does Hartry Field,²⁷ while granting that knowledge is more than justified true belief. One will then need a naturalistic account of the fourth condition for knowledge and it is difficult to retain naturalistic purity for such a condition. I do not have the space here to give this issue its due, but let me briefly illustrate the problem. The approaches that do not employ a theory of evidence—approaches such as the relevant alternatives approach and counterfactual approaches in terms of sensitivity and safety—have difficulty avoiding an appeal to the evidentially-laden concept of sameness of total epistemic situation. Relevant alternatives theories normally have to appeal to the concept of close worlds to clarify what makes an alternative relevant, making such theories equivalent to ones which emphasize safety or sensitivity (where these are clarified respectively by the claims *if the person were to hold the belief, it would be true* and *if the*

claim were false, the person wouldn't believe it). Counterfactual theories of safety and sensitivity appear to need some notion of sameness of total epistemic situation, however. Possible angels might arrange things so that the sensitivity and safety conditions are satisfied only because new information is given to the cognizer in close worlds, or this difficulty might arise through unusual features due to Mother Nature. In such cases, it is tempting to buttress one's theory by restricting the counterfactuals to require sameness of total epistemic situation. Such a theory, however, has only a thinly-veiled appeal to the concept of evidence: what it is to "have" information and "ignore" it is to fail to respond to its evidential force.

There is further problem for versions of attitudinalism that attempt a cognitivist account of the fourth condition, besides the above one of naturalistic purity, one concerning the difficulty of accounting for the value of knowledge over that of its subparts. As Timothy Williamson puts the point, "[The importance of knowledge] would be hard to understand if the concept *knows* were the more or less ad hoc spral that analyses have had to become; why should we care so much about *that?*"²⁸ It is instructive that the two problems that lead to an interest in attitudinalism re-emerge if one's attitudinalism does not cover the entirety of the gap between knowledge and true belief.

So suppose the naturalist agrees and adopts attitudinalism about the entirety of what must be added to true belief to yield knowledge. Simple versions of attitudinalism have a difficulty we might label "the Spock problem." Purely cognitive beings are possible, even if their survival as a species or subspecies is in jeopardy because if it. Such beings can both have knowledge and ascribe it to others, contrary to the claim of simple attitudinalism that such ascriptions amount to cheering for a particular state of mind.

Still, simple attitudinalism is not the only variety of that viewpoint, and it is not possible here to consider all possible varieties of the position in order to assess its prospects for inclusion in the naturalistic framework. Instead, I want to consider the version adopted by Hartry Field, inasmuch as it depends on what I and many others take to be the most plausible contemporary version of attitudinalism in ethics, that of Allan Gibbard.²⁹ Gibbard's account is supposed to apply equally to moral and epistemic norms, and he argues that, given an account of natural, Darwinian representation, there is no need to posit normative facts to explain what we are doing when we express normative judgements. Instead, our behavior can be explained solely in terms of the evolutionary value of coordination between minds regarding which norms to accept and which not to accept. This approach undergirds Field's attitudinalism about justification in epistemology.

The crucial issue for this extension of expressivism outside the domain of morality and action and into the realm of epistemology is this: any defense of the view must appeal to *arguments* and the adequacy of various proposed *explanations*, and these arguments and explanations rely on *epistemic norms*. For example, if I infer that evolutionary theory is better than some alternative because it posits fewer theoretical entities, I rely on the norm that in the search for truth Occam's Razor should be followed. For another example, if I deduce some conclusion by the canons of first-order theory, I rely on the norm that these canons are appropriate in extending our knowledge. And if I reason in accord with Bayes' theorem on some probabilistic matter, I rely on the principle that this theorem is a suitable guide in determining what to believe. The problem these points raise is that, on the face of it, if the norms are not true, the arguments and explanations fail to provide sufficient epistemic grounds for endorsing their conclusion.

Gibbard remains unmoved, however:

My explanations were of course guided by norms—epistemic norms. Why, say, did no basic tendency toward perfection figure in the explanations [of the beliefs of physicists regarding electrons] I gave? No such thing should be posited, I assumed, when observed patterns can be explained just as well without it. This is a normative judgment, and it and others like it guided me. . . . The norms that guide explanation, though, are not themselves parts of the explanation. I did not suggest that we developed our normative capacities because basic tendencies to perfection should not be posted gratuitously. Epistemic norms tell us what constitutes a good explanation, but that does not make them part of that explanation.³⁰

Of course, Gibbard is right that the norms that guide reasoning are not themselves part of that very reasoning. Norms are not *constituents* of explanations, but that is not the appropriate level of concern here. Explanations have *presuppositions* as well as constituents, and the adequacy of Gibbard's explanation depends on the truth of Occam's Razor in the way distinctive of at least some of these.

What is a presupposition? I offer no general theory of them, but we can say this much: the class of presuppositions is a subset of the class of implications of a statement or set of statements. One quite clear example is the following: the validity of modus ponens presupposes the truth of the corresponding conditional to modus ponens. In the quote above, Gibbard admits that he *assumed* the truth of Occam's Razor and that it guided his reasoning. Those claims are true, but more is true as well. It is also true that Gibbard's argument would be

inadequate if Occam's Razor (properly formulated) were not true. That is, the quality of his arguments depends crucially on Occam's Razor, in just the way that the validity of modus ponens depends on the alethic status of its corresponding conditional. In the language I am employing, his arguments presuppose the truth of Occam's Razor. If Occam's Razor is false, Gibbard had better look around for a better defense of his attitudinalist views, for the present argument supports his position only if that methodological principle is true.

One can get distracted here by the metaphysically tendentious language of facts, wondering about the distinction between facts and values, but that is not the issue at all. The question is not whether Occam's Razor is a fact, but rather whether it is a presupposition of certain explanations. It's being such requires that it be semantically evaluable, and whatever else it demands in terms of the existence of facts and the like depends on results in the theory of truth. The simple point is that arguments and explanations presuppose the truth of epistemic norms, and if the norms themselves are given non-alethic, attitudinal interpretations, then the explanations and arguments are simply defective in virtue of the fact that their presuppositions are not true.

One might think Gibbard has something of an answer here, in the following way. It is a well-known problem for expressivism how to account for reasoning involving normative conditionals (known as the Geach problem), and Gibbard employs the notion of a normative-factual world (an NF-world) to handle this problem, and one might think the same apparatus can help here.³¹

Here's how Gibbard approaches the Geach problem. We first define an NF-world: it is the opinion-set of a fully opinionated person, one who has a complete and consistent set of opinions containing for every factual claim p , either p or $\sim p$ and for every normative claim r , either r or $\sim r$. The appeal to consistency here cannot be understood in usual alethic terms, for the expressivist doesn't think that normative claims are strictly true or false. Instead, consistency for normative judgements is to be understood something like the following: having a set of opinions that is not dilemma-inducing, i.e., there is no option such that one's opinions both prohibit and permit that option.³² We will see that this need for a non-alethic conception of consistency causes problems, but before we can appreciate the problems, we need to explain the relationship between the attitudes of ordinary cognizers and this technical notion of an NF-world. We define the opinions of less-than-fully-opinionated in terms of those of fully opinionated ones: such opinions are identified with the way the content of any possible mental state is the disjunction of all the NF-worlds that might turn a partially opinionated person into a fully opinionated person. Once we have this apparatus

of NF-worlds in place, we can define validity in the usual way, except that we use NF-worlds instead of possible worlds and thereby avoid any need to appeal to the truth or falsity of normative elements in embedded contexts.

The hope of the attitudinalist is that this apparatus of NF-worlds be used to clarify the way in which an argument might presuppose a norm. Take, for example, an argument that presupposes Occam's Razor. The expressivist approach to such a case requires explaining this notion of entailment in terms of the non-logical notion of consistency outlined above. Instead of claiming that the argument can't be a good one without the norm being true, one will have to argue that anyone who accepts the argument must also accept Occam's Razor, that an inconsistency in attitude is generated in the opinions of a person by accepting the argument and rejecting Occam's Razor. That claim, however, is simply false. As Quine has taught us, nearly any combination of beliefs can be maintained if one is willing to make enough adjustments elsewhere in the system of beliefs. For example, one might accept an argument that presupposes Occam's Razor while at the same time denying that norm, claiming that one's principle is to accept arguments like this one that originally occur to one on Monday morning (we assume that the time the argument first occurred to one is Monday morning). Suitable adjustments elsewhere in the belief-system can be made to avoid any inconsistency in attitude. Gibbard will be hard-pressed to guarantee that enough questioning will unearth some hidden inconsistency, and even if some fully opinionated people would be found inconsistent, others surely would not. Little familiarity is needed with mental illness to realize that obvious entailments can be consistently denied if one is willing to change what is ordinarily accepted in a radical enough fashion.

The proper conclusion to draw is that the presuppositional relationship between norms and arguments is a different issue from the Geach problem, and one that cannot be solved by the resources Gibbard uses to address the Geach problem. What makes the problems different is that the norms in the Geach problem are explicit in reasoning, whereas presupposed norms are not. They can thus be rejected by the reasoner without inducing any inconsistency in attitude, even though they are logically inconsistent with the adequacy of the argument.

Those inclined toward attitudinalism might take the above discussion as a reason to take back any acceptance of the idea that arguments presuppose norms. They might think that the Quinean point that nearly anything can be embedded in some consistent system or other is a reason for thinking that the idea that an argument can presuppose a norm is simply mistaken. Such a response is extreme. What can be accepted, even rationally accepted, is simply not an infallible guide to entailment relations, here or elsewhere.

There are a number of illuminating points of view on this attitudinalist mistake. From a Fregean point of view, the mistake is that psychologizing some logical relationships is essential to the attitudinalist program, and there is no more reason to think that such limited psychologism is any better than global psychologism. Another perspective on the mistake is broadly Chisholmian. A large part of the history of the last century of epistemology involves the search for true epistemic principles, and some of these principles link non-normative features of the world (e.g., sensory appearances, coherence, etc.) with certain levels of epistemic acceptability. The above problem shows that expressivism cannot countenance this search, for they cannot explain how any of them could be true. In a way, this result may not be that surprising: such principles are fodder for supervenience claims about how the truth of normative claims depends on non-normative features of the world. Perhaps the lesson is that expressivism is far from an alternative explanation of the same data that non-expressivists try to explain: it is, instead, a position forced to deny the data itself.

There is one other illuminating point of view from which to view the above objection to expressivism. If we put this argument in Quinean form, it has especially devastating consequences for the naturalists' respect for science. In Quinean form, what I am arguing for is that epistemic norms have to be alethically evaluable since our best scientific theories require it—indeed, we cannot explain the concept of a *best* theory without presupposing the truth of some such norms. To turn expressivist about these norms lands one in the same turnip field as the sociobiologists who assert that we have no need for the concept of truth in an understanding of the development of science, failing to take into account that no such viewpoint can be defended by appeal to best explanation or adequate evidence of any sort. Good arguments presuppose logical and epistemic norms, and the concept of presupposition itself involves the concept of semantic evaluation.

Conclusion

Naturalism is therefore in serious jeopardy of being unable to account for or explain the value of knowledge. Non-cognitivism about epistemic norms undermines every attempt to defend any position, and cognitivist theories have very few good answers to the swamping problem. To the extent that they are successful in avoiding this problem, they do not maintain naturalistic purity, and thus fail to present an acceptable naturalistic

account of the value of knowledge. The conclusion to draw, then, is it is very hard to see how naturalism is compatible with the value of knowledge.

Notes

i[i] Richard Feldman, "We Are All Naturalists Now," American Philosophical Association, Minneapolis, May 2001.

ii[ii] Scott Sturgeon, "Comments" at the Rutgers Epistemology Conference, April 19, 2002.

iii[iii] Ibid.

iv[iv] Roderick Chisholm, *Theory of Knowledge*, 2nd edition (Prentice-Hall, 1977).

v[v] Richard Feldman, "Naturalized Epistemology," *Stanford Encyclopedia of Philosophy*,
<http://plato.stanford.edu/entries/epistemology-naturalized/>

vi[vi] Bas van Fraassen, "Science, Materialism, and False Consciousness", pp. 149-181 in Jonathan Kvanvig (ed.) *Warrant in Contemporary Epistemology: Essays in Honor of Alvin Plantinga's Theory of Knowledge*. Rowman Littlefield, 1996.

vii[vii] One might wish, at this point, to distinguish between naturalism so construed and physicalism, the view that everything in the universe is either physical or physically realized. Naturalism as construed here is logically distinct from physicalism. There is no a priori reason to assume that the methods of science will discover only what is physical, and there is no a priori reason that everything physical is within epistemic reach of the methods of science. But the distinction between the two is not center stage here, so I ignore the distinction in the text.

viii[viii] But cf. Bertrand Russell, who claimed "The law of causality . . . is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm." "On the Notion of Cause," *Proceedings of the Aristotelian Society*, xii (1912-13), pp. 1-26.

ix[ix] Alvin Plantinga, *Warrant and Proper Function*, (Oxford, 2001), chapter 11.

x[x] Alvin Goldman, "What is Justified Belief?" in George Pappas, ed., *Knowledge and Justification*, (Dordrecht, 1979), pp. 1-25.

xi[xi] *Ibid.*, p. 20.

xii[xii] Plantinga, *op. cit.*, p. 46: "So the view I propose is a radical naturalism: striking the naturalistic pose is all the rage these days, and it's a great pleasure to be able to join the fun."

13 Plantinga, *op. cit.*, see especially his views on the concept of a function.

14 For more on different kinds of defeaters, see John L. Pollock, *Contemporary Theories of Knowledge*, (Rowman and Littlefield, 1986).

15 The generality problem was first presented in Richard Feldman and Earl Conee, "Evidentialism," *Philosophical Studies* 48 (1985): 15-34. The problem, as they present it, is that if the processes are individuated too coarsely, then, e.g., all perceptual beliefs are justified or none of them are (because they are all produced by the same process). If the processes are individuated too finely, then only the belief in question will be produced by that process, yielding the result that the process is reliable, and hence the belief justified, if and only if the belief is true. The problem thus issues a challenge to reliabilists to find the right level of generality for their theory.

16 Ernest Sosa, "Intellectual Virtue in Perspective," in *Knowledge in Perspective*, (Cambridge, 1991), p. 284.

17 *Ibid.*, p. 287.

18 *Ibid.*, pp. 289ff.

19 John Greco, *Putting Skeptics in Their Place*, (Cambridge, 2000), p. 218.

20 Ibid., p. 191.

21 See Laurence Bonjour, *The Structure of Empirical Knowledge*, (Harvard, 1985), pp. 38-45. I'm thinking here especially of the case of Norman, p. 41.

22 See, e.g., "The Basic Notion of Justification," Christopher Menzel, co-author, *Philosophical Studies* 59 (1990), pp. 235-261; "Plantinga's Proper Function Theory of Warrant," *Warrant in Contemporary Epistemology*, J.L. Kvanvig, ed., (Savage, MD: Rowman & Littlefield, 1996), pp. 281-306; "Zagzebski on Justification," *Philosophy and Phenomenological Research* 60 (2000), pp. 191-196; and *The Intellectual Virtues and the Life of the Mind: On the Place of the Virtues in Contemporary Epistemology*, (Savage, MD: Rowman & Littlefield, 1992), chapter 5.

23 Plato, *Meno*, 97c-d.

24 See, e.g., Roderick Chisholm, *Theory of Knowledge*, 3rd edition (Prentice-Hall, 1989).

25 See, e.g., Wayne Riggs, "Understanding Virtue and the Virtue of Understanding," John Greco, "Knowledge as Credit for True Belief," and Ernest Sosa, "The Place of Truth in Epistemology," all forthcoming in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Michael DePaul & Linda Zagzebski, eds. (Oxford: Oxford University Press).

26 I will not attempt in the text any clarification or taxonomy concerning the varieties of non-cognitivism, and will vary my terminology regularly between terms that ought to be distinguished. Attitudinalism ought to be a view that identifies the non-cognitive element in terms of the *attitudes* of people holding such beliefs, whereas expressivism ought to be a theory about what is involved in certain types of speech acts. These issues are important, but not in the present context, so I will ignore them in the text.

27 Hartry Field, "The A Prioricity of Logic," *Proceedings of the Aristotelian Society* 96 (1996), pp. 359-379; "Epistemological Nonfactualism and the A Prioricity of Logic," *Philosophical Studies* 92 (1998), pp. 1-24.

28 Timothy Williamson, *Knowledge and Its Limits*, (Oxford 2000), p. 31.

29 Allan Gibbard, *Wise Choices, Apt Feelings*, (Harvard, 1990).

30 Gibbard, *op. cit.*, p. 122.

31 I thank Matt McGrath and Jeremy for this suggestion, and for other very helpful comments on an earlier draft.

32 Thanks to Robert Johnson for this idea.