

Epistemological Externalism and the Project of Traditional Epistemology

Contemporary philosophers still haven't come to terms with the project of traditional epistemological reflection on our knowledge of the world. In this project we bracket all our beliefs about the world, foregoing our ordinary readiness to utilize them in our intellectual endeavors, and from that vantage point attempt to account for our knowledge of the world.¹ In one familiar incarnation of this project, for instance, we ask how one could “get from” one’s sensory experiences, described in a way which is noncommittal as to how the world is, to knowledge about the world. In an even more common manifestation of the standpoint involved in this project, it is assumed that one’s ultimate evidence for one’s beliefs about the world cannot, in every case, involve other considerations about the world. Many philosophers have argued that this project inevitably results in skepticism; others have been less pessimistic. But is the project even intellectually sensible? Is it reasonable to seek to account for our knowledge of the world in this way? A wide tradition of Twentieth-Century thought, including Quine, Wittgenstein, and J. L. Austin, urged that it is not. This paper is part of a detailed defense of that answer.²

My argument will focus upon the relation between contemporary epistemological externalism and the traditional project. It is sometimes felt that externalism has merely

changed the subject,³ capitulated to skepticism,⁴ or at best offered an unsatisfactory answer to the traditional question.⁵ As I will argue, however, externalism can provide us with principled grounds for rejecting the traditional project. My point is not that externalism is inconsistent with certain theses or requirements involved in the traditional project. Rather, it is that externalism undercuts the most plausible intellectual motivations for engaging in the traditional project in the first place.

To show this, I will bring out the consequences of a thesis which I call *Minimal Externalism*:

It is a necessary condition for having knowledge (or for having justified belief) about the world outside of one's mind that certain conditions, not concerning the truth-value of the belief, are met in the world outside of one's mind. (These conditions may be only negative, requiring that certain things not be the case.)⁶

I will argue, roughly, that if this thesis is true, then the traditional project will not enable us to account for our knowledge of the world, nor will it have any relevance for whether we have knowledge or justified beliefs about the world. It therefore should be abandoned. As we will see, this result stands even if one holds that propositional knowledge (at least in the case of mature human beings) ordinarily requires justified belief, and that (subject to the same proviso) being justified ordinarily requires the ability to offer adequate reasons in defense of one's belief. Consequently, a decisive break with

the epistemological tradition can retain much of the traditional conceptions of knowledge and justified belief.

A familiar criticism of the traditional project is that by raising a global question about our knowledge of the world as a whole, it violates certain constraints on the activities of justifying or evaluating beliefs. Questions of epistemic justification and evaluation, it is said, are a *local* matter, arising within specific contexts in which the vast majority of beliefs about the world are not in question and may be freely drawn upon.⁷ A common rejoinder is that this point, while true enough as a description of how things usually go, is no criticism at all: it merely shows that traditional epistemological questions are unusual questions, not the sort that would be raised in any ordinary practical or investigative setting.⁸ One main goal of my argument is to show that this reply fails. When framed in terms of minimal externalism, the localist reply to the traditional epistemological project has real bite.⁹

My discussion will proceed as follows. The first section articulates the most plausible motivations for the traditional project. The following two sections argue that minimal externalism can undermine these motivations. My conclusion, then, is a conditional one: if minimal externalism (or an appropriate version of it) is correct, then we have no compelling reason to engage in the traditional project and so can reasonably reject it.

Moving beyond this conditional conclusion requires a full-dress defense of minimal externalism, a task which I will not undertake here. However, minimal externalism is almost universally accepted with regard to empirical knowledge; what Gettier taught us is precisely that external-world conditions other than the truth value of one's belief determine whether or not one has knowledge of the world.¹⁰ There is much less agreement about whether minimal externalism is true with regard to justification. Here, roughly and dogmatically, is how I would motivate it. Justified belief (and, I believe, knowledge) requires that one have good epistemic reasons for believing as one does. Good epistemic reasons are considerations which tell in favor of the truth of one's belief. The relevant notion of "telling in favor of" is, I believe, an objective notion relating to the objective likelihood that one's belief is true; roughly, on my view, it is the notion of *being a good indicator*. What is a good indicator of the truth of a particular claim about the world depends upon the world: upon what is the case, what the relevant things are like, and how matters tend to go. Accordingly, in many if not in all cases, whether certain considerations are good reasons for a particular belief about the world depends upon matters in the world outside of the person's mind. However, nothing I say here will require accepting the details of this view. My concern is simply the relation between such views and the traditional project.

Three additional points should be noted. First, my topic is not historical. I will not speculate about any particular philosopher's actual motives for undertaking the

traditional project, but will instead consider how we might best motivate that project for us, now. Second, I will not consider motivations arising from the desire for absolute certainty or infallibility. Such hankerings, tempting though they may be, no longer provide good intellectual motivation for epistemological projects. Finally, my primary concern is not skepticism. Rather, it is the motivation of, and prospects for, an epistemological project which the skeptic and anti-skeptic arguably share.

I. Motivating the traditional project

The traditional epistemological project aims at developing a fully general theory or account of our knowledge of the world. This project is guided by a fundamental ground-rule, which I will call "the Cartesian Constraint:"

A complete philosophical theory of our knowledge of the world must not ineliminably appeal to any considerations about the world.

What intellectual motivation could one have for undertaking a theoretical project structured by this constraint?

It is frequently assumed that the Cartesian Constraint is motivated by disputed "internalist" principles about knowledge or epistemic justification. One principle which is often thought relevant is an *accessibility requirement*:

In order to know or justifiably believe something, one must be in a position to determine through introspection and *a priori* reflection alone that one knows or justifiably believes it.¹¹

This requirement does motivate a project very much like the traditional one. If the accessibility requirement is correct and one has knowledge or justified beliefs about the world, then one should, in principle, be able to account for one's knowledge or justified beliefs about the world merely in terms of the resources of introspection and *a priori* reflection. One might therefore test whether one has knowledge or justified beliefs about the world by seeing whether one can do what the accessibility requirement says one must be able to do: determine whether one has knowledge or justified belief about the world without ineliminably relying upon any considerations about the world. Still, this requirement should not be our focus, for two reasons. First, the accessibility requirement is not necessary in order to motivate a project structured by the Cartesian Constraint. As we will see shortly, a plausible and less tendentious motivation is available. Second, the accessibility requirement is only as plausible as the argument for it, and, so far as I can see, no compelling argument has ever been put forward.¹²

A second principle which may be thought relevant here is a *reasons requirement*, a principle linking knowledge or justified belief with the ability to provide adequate reasons for believing as one does.¹³ However, this principle, too, is not necessary to motivate a project structured by the Cartesian Constraint (though one of the motivations

which I will discuss at length does require it). Moreover, combining this principle with certain forms of minimal externalism can provide grounds for rejecting the traditional project, as I will argue in section 3.¹⁴

Plausible motivation for the traditional project can be found by considering a suggestion from Barry Stroud. According to Stroud, the traditionalist desires to answer the question, "How does anyone know anything at all about the world?" This question invites an account which would explain *all* of our knowledge of the world. This very generality, Stroud claims, generates the Cartesian Constraint.

The demand for a completely general understanding of knowledge in a certain domain requires that we see ourselves at the outset as not knowing anything in that domain and then coming to have such knowledge on the basis of some independent and in that sense prior knowledge or experience...¹⁵

This proposal is on the right track. However, we must distinguish several different projects which could be described as aiming to provide a "general explanation of our knowledge of the world." Not all of them yield the Cartesian Constraint.

The desire for a fully general account of our knowledge imposes three requirements. First, the account must apply to *each* of the beliefs in question. Second, it must be a *general* account, in that it must explain all the cases in a single, uniform way.¹⁶ Finally, in order to provide what you take to be a satisfactory explanation of something,

you must at least take yourself to know the truth of the considerations appealed to in the explanation. As we will see, this last point is of considerable importance.

The simplest general explanatory project attempts to state what it is in virtue of which we possess knowledge of the world.¹⁷ This project has two parts: a statement of the conditions under which a person has knowledge and a description of the particular way(s) in which we meet these conditions. Taken together, these two parts would explain why what we have is *knowledge*.

This project does not generate the Cartesian Constraint. To see this, suppose that one holds that the conditions for having knowledge involve factors in the world beyond the individual's mind. One will consequently appeal to considerations about the world, considerations which one thinks one knows to be true, in order to explain how what one has amounts to knowledge. Nonetheless, this explanation will be fully general, provided that it applies to each instance of knowledge and (at an appropriate level of characterization) does so in a uniform way. None of one's knowledge of the world will remain unexplained.

It might seem that by appealing to things which one takes oneself to know about the world, the explanation loses a further desirable kind of generality. We might want our account to reveal *what knowledge of the world is*, what it consists in, without essentially utilizing the concept of knowledge. If so, then we should not allow the statement that one knows some particular thing about the world to figure as an ultimate

constituent of the explanation of why what one has amounts to knowledge of the world. However, this restriction does not preclude considerations about the world from appearing in the explanation. There is a difference between an explanation which uses claims about one's knowledge and an explanation which merely uses claims which one takes oneself to know. In the latter case, it is p which does the explanatory work — not the proposition that one knows that p .¹⁸ Likewise, it should not be thought that this explanatory procedure would beg the question, since the question at issue in this explanatory project is not whether one knows things about the world, but merely what it is in virtue of which one knows them.

One version of this project deserves particular notice. Assuming that knowledge requires evidential justification, some epistemologists such as Chisholm and Roderick Firth sought to explain our knowledge of the world by “reconstructing” our justification for believing anything about the world. Firth described the project as follows.

To reconstruct our knowledge is simply to order the propositions that we know in the form of an argument [schema] that exhibits the structure of the justifying relations that hold among them. ... [The reconstruction] will thus show, in a schematic way, the extent to which certain parts of our empirical knowledge are dependent for their justification or warrant on other parts, and it will formulate the principles of non-deductive inference ... in virtue of which these justifying relations obtain.¹⁹

By itself, the demand for such a reconstruction does not generate the Cartesian Constraint. For instance, it is perfectly consistent with this project to hold that at least some of the relevant principles of non-deductive inference are contingent and only known

empirically. On such a view, considerations which one takes oneself to know about the world would be an ineliminable part of a reconstruction of one's justification for one's beliefs about the world. Nonetheless, the explanation would be fully general, since it would explain what it is in virtue of which any of one's beliefs about the world constitute knowledge. This point is obscured by the fact that traditional epistemologists such as Chisholm and Firth generally assumed the relevant principles of non-deductive inference to be necessary and knowable a priori. But this assumption is not generated simply by the attempt to explain with complete generality what it is in virtue of which we possess knowledge of the world.²⁰

A second type of general explanatory project seeks a causal or "genetic" account describing the events, processes, or mechanisms by which we come to have knowledge about the world. It is sometimes claimed that when one applies this project to oneself, one can no longer view oneself as knowing the things which one must appeal to in order to provide the explanation.²¹ However, nothing in the simple causal or "genetic" explanatory question forces one to detach oneself from one's own beliefs in this way. Like the previous question, it does not call into question any of one's knowledge. Accordingly, one could answer it by appealing to things which one takes oneself to know about the world. Consider, for illustration, a crude reliabilist theory according to which knowledge is true belief formed through a psychological process or mechanism which yields true belief in an appropriately high percentage of cases. A theorist who holds this

view might study the mechanisms through which people arrive at their beliefs. He might find that they are in fact reliable. And so by describing these mechanisms and noting that they are reliable, he might explain how people, including himself, gain knowledge. He can thus use what he knows of the world to provide a satisfying general “genetic” explanation of his knowledge of the world.²²

There are, however, two general explanatory projects which *do* give rise to the Cartesian Constraint.

The first arises from a certain way of generalizing the question, “How do you know?” In many ordinary conversational contexts, this question functions as a request for a justification of your belief or assertion. In response, you are supposed to offer reasons which adequately support your belief’s truth without begging the question. (If you are unable to do so, then — assuming no special considerations apply — your interlocutors will deny that you know the thing in question.) Justifying explanations can be requested regarding classes of beliefs as well as individual beliefs. Accordingly, it seems that one could intelligibly ask oneself such questions as, “How do I know anything at all about the world outside my own mind?” intending thereby to request a justification for those beliefs. To answer this question, one would have to provide a general argument justifying holding beliefs about the world at all. It should be clear that one cannot satisfactorily answer *this* question by appealing to or presupposing considerations about

the world. To do so would beg the question. (In what follows, I will call this project the “general justificatory project”.)

The second general explanatory project which yields the Cartesian Constraint arises from a certain kind of reflective curiosity about the state of our knowledge.²³ We sometimes ask whether we actually know certain things which we think we know. We answer such questions by considering whether our evidence or the way in which we formed the belief(s) are adequate for knowledge. We thus establish that we possess the knowledge in question by explaining how we know. I will call such explanations “vindicating explanations”. It should be emphasized that the request for such an explanation does not require the thought that one’s possession of knowledge requires one to be able to provide such an explanation.

One can wonder, in a reflective moment, whether one really has any knowledge of a certain kind just as well as whether one really knows some particular thing. So curiosity could lead one to ask, “Do I really know *any* of the things which I take myself to know about the world?” To answer this question, one could not appeal to anything which one takes oneself to know about the world. For a satisfactory vindicating explanation cannot make use of the very claims one’s knowledge of which is being explained; one cannot answer the question, “Do I really know that p?”, by appealing to p itself. The best explanation of this fact involves the pragmatics of assertion or explanation. Appealing to p in this particular explanatory context presupposes, or

commits one to, the claim that one knows that *p*.²⁴ But whether one really knows that *p* is precisely what is at issue.

The search for a vindicating explanation of our knowledge of the world is an instance of an attitude of reflective distancing which we can take towards our beliefs more generally. We can regard a given belief as a mere fact about our psychology, bracket our commitment to the truth of its content, and ask, "Is this belief really true? Are things actually as I take them to be?" From this perspective, it will appear insufficient to answer the question merely by saying, "Well, things *are* after all that way;" that reply would acquiesce in the belief, not vindicate it. The traditional epistemological project can thus arise when we take this stance towards our beliefs about the state of our knowledge of the world. Understood in these terms, the project's specifically epistemological commitments are quite minimal. In particular, it requires neither accessibility internalism nor the reasons requirement.²⁵

These considerations clarify the state of play between minimal externalism and the traditional project. It has been charged that externalism cannot provide a satisfying general account of our knowledge of the world because the externalist must appeal to claims which she takes herself to know about the world as part of the account.²⁶ A standard response is that externalism yields a satisfying general explanation if we (1) grant that one can't satisfactorily explain one's knowledge unless one knows or justifiably believes the claims which provide the explanation, but (2) deny that one must

know or even believe that one meets this requirement.²⁷ This reply is not satisfying. For one thing, neither the general justificatory question nor the general vindicatory question requires second-order epistemic requirements for its motivation. Moreover, it is clearly question-begging to appeal to considerations about the world in response to the general justificatory project. And unless a lot more is said about the pragmatics of assertion and explanation, it is not plausible that one can provide a satisfying explanation of one's knowledge by appealing to considerations which one doesn't even take oneself to know or justifiably believe to be true. As we have seen, however, this last point does not prevent an externalist from successfully undertaking the first two general explanatory projects. So the initial objection is not quite right either; only the general justificatory and vindicatory questions disallow considerations about the world. The crucial question, then, is whether minimal externalism provides good grounds for rejecting these latter two motivations for the traditional project.

II. Minimal Externalism and the Project of Global Vindication

A fully successful criticism of the general vindicatory project would show that it is an *intellectually irrelevant* exercise. This is precisely what the minimal externalist is in a position to charge. If minimal externalism is true, then (1) one cannot develop a general vindicatory explanation of one's knowledge of the world, (2) this fact has no significance for, and does not reveal anything about, the actual epistemic status of one's beliefs about

the world, and (3) there is another way of determining whether or not one really knows things about the world. I will treat each of these points in turn.²⁸

Regarding the first point. The vindicatory project asks, about a certain class of claims *W*, whether we really know any of them. If one accepts minimal externalism, then one holds that some condition(s) *CW*, concerning the world outside one's mind, must be met in order for one to know things in *W*. So in order to explain, in general terms, how one meets the conditions for knowing things in *W*, one must take there to be some circumstances in virtue of which conditions *CW* are satisfied. However, claims such as "There are circumstances in virtue of which *CW* are satisfied" are themselves claims in *W*. As we've noted, in order to appeal to such claims in the course of one's explanation, one must take oneself to know them to be true. But whether or not one knows them is part of what one's investigation is meant to establish, since one is attempting to determine whether one knows *anything* in *W*. So the initial question, "Do I really know anything about the world at all?", prevents one from appealing to the relevant claims in the course of one's explanation. Thus, if minimal externalism is true, one cannot even begin to develop a fully general vindicatory account of one's knowledge of the world. This is not to say that a minimal externalist is guaranteed to reach a negative verdict. Rather, she lacks a basis for answering the question at all.²⁹

It would be incorrect to conclude that there is a well-motivated intellectual project which the minimal externalist cannot complete. For the above considerations provide the

minimal externalist with a principled criticism of the general vindicatory project: that project removes the background framework of claims about the world which makes it possible to determine whether you know any particular thing. This criticism brings us to the second point which the minimal externalist should make about the general vindicatory project. If minimal externalism is true, then our inability to succeed in the general vindicatory project arises simply from the interplay between the requirements of the general vindicatory project, the consequences of minimal externalism, and the particular class of beliefs chosen for reflective scrutiny. It is completely unrelated to the actual epistemic status of our beliefs about the world, and it therefore does not indicate anything at all about whether or not we possess knowledge.³⁰ So if the general vindicatory project is to reveal something about the actual state of our knowledge, minimal externalism must be false. But, the minimal externalist claims, minimal externalism is true. Consequently, the minimal externalist can object as follows: if you want to know whether you really have knowledge about the world, *this isn't the way to find out*. The general vindicatory project simply blinds you to whatever knowledge you in fact possess.

This brings us to the minimal externalist's third point. The minimal externalist will insist that to find out whether one really knows things, one should proceed piecemeal, answering the question for particular beliefs (or classes of beliefs more specific than the class of beliefs about the world) on the basis of background beliefs

about the world which remain available once the relevant beliefs have been bracketed.

To determine whether we *really* know X or Y about the world, the minimal externalist asserts, we need not first determine whether we really know *anything at all* about the world. It should not be felt that this reply simply changes the subject. The minimal externalist's piecemeal project and the general vindicatory project concern the same subject matter: both aim to determine whether some of our ordinary knowledge claims about the world are actually true. The general vindicatory project simply aims to do so by considering all of them all at once, while the minimal externalist urges an alternative approach.

Is this alternative approach intellectually satisfying? It is true that on the minimal externalist's method the conclusions we draw about our knowledge will only be as reliable as the background beliefs upon which we have relied. But this much is true of any intellectual endeavor, and we certainly shouldn't be dissatisfied with the minimal externalist's position simply because it can't yield infallible conclusions about the state of our knowledge. If we have no reason to think that our belief system is largely incorrect, then the minimal externalist's method is the most intellectually responsible way to attempt to find out whether we really know things which we think we know. (If one had what appeared to be a good reason to doubt that one knows anything at all about the world, then the general vindicatory question would be quite appropriate, even given

minimal externalism, and our inability to answer it affirmatively might well be intellectually disappointing. Fortunately, however, this is not our situation.)

The minimal externalist can therefore charge that no good intellectual purpose is served by posing the question, “Do I really know anything at all about the world?” The charge is not that this question is meaningless or somehow less than fully intelligible. Rather, it is that the desire to answer it is misguided in much the way that it would be misguided to desire, merely out of curiosity, to visually determine whether there is furniture in a completely dark room without turning on the lights. One could desire such a thing, but it would be foolish to do so; the resulting research project is patently pointless. It would be better to turn on the lights.³¹

III. Minimal Externalism and the General Justificatory Project

I turn now to the motivation for the Cartesian Constraint provided by the general justificatory project. The general justificatory project assumes that knowledge requires justification, and it demands good reasons which can be offered in justification for beliefs about the world without presupposing or depending upon any beliefs about the world. These reasons are to be *general*, in that they should support the proposition that one’s beliefs about the world are for the most part true.³² (They may also provide a general justificatory schema, instantiated in particular ways for particular beliefs about the world.) And the project assumes that unless we are in a position to provide such reasons,

no belief about the world is justified — regardless of whether we possess good reasons of some other sort for thinking it to be true.

This demand can be motivated by a model from our ordinary practice. Suppose that I defend a claim in physics by appealing to other claims in physics as support. Someone might ask, "How do you know anything at all about physics? What reason do you have for believing any of this stuff at all?" Here the interlocutor requests good reasons, from outside the domain, which can be offered as good reasons even when I have bracketed every belief within the domain. (I might say, for instance, "I took some physics in college and have kept up with recent popular books in the field.") And if I wasn't in a position to offer some such defense, my interlocutor would regard me as not being justified, all else equal. This strongly suggests that the justifications by which I would ordinarily defend particular claims in physics are part of a larger justificatory story: they justify me in holding particular beliefs in the domain only if I also possess good reasons for holding any beliefs in the domain at all. The suggestion behind the general justificatory project is that the justifications we ordinarily provide for our beliefs about the world are in the same boat.

Two questions need to be considered. First, is it possible to provide the sort of justification this project demands? Second, must we be able to provide such a justification in order to have knowledge or justified beliefs about the world? In this section I will argue (roughly) that if certain forms of minimal externalism are accepted,

then one can reasonably answer both questions negatively. Consequently, a minimal externalist can reasonably reject the general justificatory project.³³

My argument will also involve a subsidiary concern. The general justificatory project presupposes a version of what I earlier called a *reasons requirement*:

To be justified in believing that *p* about the world, one must be in a position to provide good reasons for believing it.³⁴

Many contemporary epistemologists reject this requirement, holding instead that one can be justified even if one is incapable of providing any good reasons for one's belief.³⁵ This view does undercut the general justificatory project, but it strikes me as too drastic.³⁶ I therefore aim to show that minimal externalism can motivate rejection of the general justificatory project even if we accept the reasons requirement.

My discussion will focus on versions of minimal externalism which accept the following two theses:

(A) Being justified requires having good reasons for one's beliefs — reasons which provide good evidence for the truth of one's beliefs, and

(B) The goodness of reasons is determined by circumstances in the external world, at least in certain crucial cases.

I will call the family of views which accept these two theses "minimal externalism about reasons."³⁷ William Alston's version of foundationalism provides a useful illustration.

According to Alston, a belief about the world is immediately justified just in case it is

based upon a perceptual or other non-belief psychological state which is an "adequate indicator" of its truth, where a given state is an adequate indicator of a belief's truth just if it makes the belief's truth highly likely *given the causal laws which actually govern our world*.³⁸ This is just one example. (I will say more shortly about the various forms minimal externalism about reasons can take.)

I begin with the question of whether it is possible to succeed in the general justificatory project if one accepts minimal externalism about reasons. Here I want to emphasize a point concerning the activity of offering reasons in defense of beliefs (which I will call "the activity of justifying"): many versions of minimal externalism about reasons have the consequence that if one accepts them, then (given certain facts about the activity of justifying) one cannot possibly provide a general justification for one's beliefs about the world.

To see why this is so, consider the structure of the activity of justifying. For any given attempt at justifying, there is a set of propositions which are "off-limits" in the following sense: one may not appeal to or rely upon them without providing some further considerations, not in the set, in their favor. To take one example, the belief being justified always belongs in this set. To appeal to or rely upon it in one's justification would beg the question. Other propositions may also be in this set, perhaps because of features of the conversational or dialectical setting, perhaps because of objective relations of evidential priority between propositions. Whether this is so, and exactly what factors

might operate to determine which propositions these are, are vexed matters which I will not tackle here. It suffices for my purposes that the proposition(s) being justified are always in this set. I will call this set the “suspension set”.

The suspension set for a given justificatory attempt affects what argumentative transitions are permissible in the course of the justification. For any given transition from a reason P to a claim Q, there is a corresponding proposition to the effect that P supports, tells in favor of, is evidence for, or is a good (epistemic) reason for believing Q. I will call this the “linking proposition”. If the linking proposition is in the suspension set, then unless one has some reasons for accepting it which are not themselves in the suspension set, one may not acceptably appeal to P to justify believing Q.³⁹ Suppose, for instance, that I am attempting to justify the belief that my garden has moles and that the suspension set contains the proposition that certain sorts of lawn damage are a good reason to believe that one’s garden has moles. In such a case I cannot acceptably appeal to the state of my lawn in defense of my belief that my garden has moles. Or, at least, I can’t do so unless I have some reason, not in the suspension set, in favor of the linking proposition.

Similarly, if the only reasons one could offer in favor of the linking proposition are themselves in the suspension set, then one may not appeal to P to justify believing Q unless one is in a position to support those reasons with some further reasons not in the suspension set. Suppose, for instance, that (1) my only reason in favor of the proposition linking lawn damage to moles is that moles cause this sort of damage when digging their

tunnels, (2) this last proposition is itself in the suspension set, and (3) I don't recognize any independent reason to believe it. In that case, too, I cannot acceptably appeal to the state of my lawn in defense of my belief that my garden has moles. Consequently, I cannot acceptably defend the belief, that moles cause this sort of damage, by appealing to an argument which depends upon a transition from considerations about lawn damage to the claim that my garden has moles.⁴⁰

Consider now how these points apply to the attempt to provide a general justification for one's beliefs about the world. As previously noted, this justificatory project places all of one's beliefs about the world "out of bounds" for the purposes of the justification. They are all in the suspension set. Now if one accepts an appropriate version of minimal externalism about reasons, then one will hold that the relevant considerations about reasons are all themselves considerations about the world or (at the least) are ultimately supportable only by dint of appeal to considerations about the world. Accordingly, the advocate of an appropriate version of minimal externalism about reasons must place all the relevant considerations about reasons in the suspension set for the attempt to provide a general justification. And this fact means that if one accepts an appropriate form of minimal externalism about reasons, then it is structurally impossible for one to provide a satisfactory general justification for one's beliefs about the world. The thought that something "internal" tells in favor of the truth of something "external" is itself in the suspension set; accordingly, the most one could say is, "If the world is a

certain way, then certain internal facts tell in favor of the general truth of my beliefs about the world; if not, then not. However, I can't legitimately rely upon a claim either way, given the structure of the justificatory situation created by my attempt to provide a general justification for my beliefs about the world, so in this justificatory situation I cannot legitimately appeal to anything in defense of the claim that my beliefs about the world are generally true."

This is not to say that the resources available at the beginning of the attempt are evidentially inadequate or that they do not reliably indicate the truth of one's beliefs. The trouble is rather that if all claims about the world are in the suspension set, then one will not be in any position to make justificatory use of whatever good evidential resources one has. Consequently, my claim is not this: if minimal externalism about reasons is true, then nothing indicates that your beliefs about the world are generally true. Rather, my claim is that if you accept minimal externalism about reasons and are clearheaded about the situation, you will see that once you ask the general justificatory question, you simply aren't in a position to justify any beliefs about the world at all. You can't so much as get started.⁴¹

According to some forms of minimal externalism about reasons, all relations of evidential support (aside from purely logical relations) are grounded in causal patterns in the world.⁴² The inability to provide a general justification follows directly from acceptance of any such view, since any evaluation of p as a reason for believing q would

then involve or depend upon considerations about what is the case and the way things tend to go in the world. It might seem that a foundationalist version of this view could avoid this result, since it could allow that the relevant beliefs about the world are ultimately warranted simply by the fact that one has certain appropriate sensory experiences. However, if there is an evidential support relation between the fact that one has certain sensory experiences and the truth of propositions about the world, it is not a purely logical relation. Consequently, if one accepts the form of minimal externalism about reasons currently under consideration, one could not acceptably appeal to one's sensory experiences as a justification for believing anything about the world unless one is permitted to rely upon certain propositions to the effect that the world is such that sensory experiences are a good indication of how things are in the world. But all such propositions are in the suspension set for the attempt to provide a general justification. An analogous point would apply to a coherentist attempt to reach conclusions about the world by appealing to the fact that one holds a comprehensive and appropriately coherent system of beliefs about the world. For these reasons, the impossibility of providing a general justification also follows from accepting the weaker view that evidential support relations are grounded in causal patterns in the world only in certain crucial cases (such as the evidential relations between sensory experience, or the coherence of belief systems, and the truth of certain propositions about the world).⁴³

The results reached so far have important consequences for our understanding of the reasons which are available for our beliefs about the world. Consider, again, what the general justificatory project seeks: reasons for beliefs about the world which can be deployed even when all beliefs about the world are in the suspension set. If one accepts an appropriate form of minimal externalism about reasons, then one must answer that there are no such reasons. One must hold that the only reasons which one can provide in support of any given belief or beliefs about the world are either claims about the world or considerations which one cannot appeal to in justification for beliefs about the world unless one is permitted to depend upon or presuppose some beliefs about the world. Consequently, one must also hold that if it is to be possible to satisfactorily engage in the activity of justifying, then it must be a piecemeal activity which permits one to draw upon background beliefs about the world in the course of offering reasons for particular target beliefs about the world. As I will put this result, the activity of justifying beliefs about the world will be *local*.

This isn't simply a point about what one must accept if one accepts minimal externalism about reasons. Suppose that an appropriate form of minimal externalism is true. And now imagine a being who knows all the relevant facts about the world and about the nature of epistemic reasons. This being sets out to provide a general justification for its beliefs about the world, seeking acceptable reasons for thinking those beliefs to be generally true. This being will conclude that there are no such reasons. If

what kinds of epistemic reasons there are is an objective and non-perspectival matter, then it follows that this being is right: if an appropriate form of minimal externalism is true, then there are no reasons which can provide an adequate general justification for one's beliefs about the world.

Even if one balks at this conclusion, the truth of minimal externalism would have important implications for our epistemological aspirations. The attempt to provide a general justification is part of a distinctively philosophical project of understanding our epistemic relation to the world. And what we have seen is that if minimal externalism about reasons is true, then no satisfactory general justification for our beliefs about the world can be imbedded within a correct theory of our epistemic relation to the world. That theory will have to grant that if we are able to provide good reasons in justification for beliefs about the world, those reasons can only be local – reasons which are correctly deployed in the activity of justifying only in settings in which some beliefs about the world are not in the suspension set.

Does this result show that if we accept the reasons requirement and minimal externalism, then we must accept that there is a legitimate justificatory requirement which we cannot meet? No. Consider the following two requirements:

Basic Justificatory Requirement (BJR): In order to be justified in

believing any p about the world, I must be able to justify believing p by

articulating good reasons for believing p .

Global Justificatory Requirement (GJR): In order for any of my beliefs about the world to be justified, I must be able to provide a general justification for my beliefs about the world.

BJR does not entail GJR. GJR requires the ability to provide reasons which are not themselves considerations about the world and which one can acceptably offer even when all considerations about the world are in the suspension set. BJR, however, can be understood in terms of the view that the activity of justifying is a local activity; for all it says, the requisite good reasons could always include other propositions about the world. Consequently, one could accept it while rejecting the Global requirement. Minimal externalism about reasons is compatible with BJR, interpreted in local terms. And if we put them together, the resulting view is in an excellent position to argue that the desire for a global justification is misguided: on this view, no such justification can be provided (because of minimal externalism about reasons), and this fact is entirely irrelevant to the justificatory status of our beliefs about the world (because, it is claimed, GJR is false).

This brings us to the second major step in the argument against the general justificatory project. Minimal externalism about reasons buttresses rejection of GJR.

My argument here appeals to the following assumption:

If we cannot reasonably see GJR as implicit in or derivable from the requirements involved in our everyday justificatory practices, then it may reasonably be rejected.

My primary motivation for this assumption is methodological; I don't see any more promising way to decide the question.⁴⁴ In appealing to our everyday justificatory practice, I do not mean our usage of the word, "justified". My appeal is rather to the overall pattern and structure of our ordinary activities of evaluating and justifying beliefs: how the activity goes, when and how we engage in it, the requirements we explicitly or implicitly insist upon, the consequences of various outcomes, etc. The relevant approach to these matters is not the purely descriptive approach of the social scientist, but rather the viewpoint of someone whose practice it is, someone who endorses or is committed to it (at least in broad outline), and what is needed is an account of what happens when we are doing it *right*. This last point does not render the appeal to our practices question-begging, since we can agree upon exemplary instances of our ordinary practice without (yet) agreeing about whether GJR is correct.⁴⁵ It may be that scrutiny of exemplary instances of our practice will reveal good reasons for thinking that we implicitly invoke or are committed to GJR, or that GJR is derivable from some clear feature of our practice. What is rejected is simply the suggestion that though GJR is correct, there is *nothing* about the justifications which we offer, the requirements which we impose, or our patterns of response to practical and conversational circumstances, which commits us to, or indicates a commitment to, GJR.

I do not claim that all correct epistemic requirements are related to our practice in this way. If minimal externalism about reasons is true, then the requirements one must

meet in order to be justified will depend in part upon what in fact tells in favor of the belief's truth, and that need not be revealed by our ordinary practice. However, GJR concerns the fundamental structure of the justifications which we must possess in order to have justified beliefs about the world; it claims that local reasons, even local reasons which actually support the truth of our beliefs, aren't enough. And my methodological assumption expresses just the minimal idea that any reasonable claim about this matter must be supportable by appeal to our everyday practice.

Let us begin, then, by recalling some pertinent facts about our practice. First, if GJR explicitly figured in our everyday practice, then there would be ordinary occasions on which people (1) attempted, or were at least asked to attempt, to justify beliefs about the world in a way that is completely free of any presuppositions about the world and (2) were treated as holding unjustified beliefs if the attempt failed. But it seems fair to say that this *never* happens in everyday life. Second, it seems extremely implausible that we ever say or do anything in ordinary life which involves an implicit invocation of GJR. We do not ordinarily understand ourselves as offering or requesting fragments of justifications which would, if fully spelled out, satisfy GJR. (That this is what we are doing is a philosophical interpretation of our ordinary practice.) Finally, if one accepts minimal externalism about reasons, then one must hold that none of one's everyday justifications is *actually* a fragment of a larger justificatory story of the sort demanded by GJR: as I have argued, the minimal externalist about reasons must conclude that no

justification could make use of reasons of the relevant sort in the relevant way. This means that if one accepts minimal externalism about reasons, then one should hold that GJR cannot be derived from the actual structure of the justifications which one regards as acceptable in ordinary life. Our ordinary practices, in conjunction with minimal externalism about reasons, thus strongly suggest that GJR is incorrect.

This argument can be resisted even if one accepts my methodological assumption and acknowledges the facts I've just listed. To do so, one must (1) explain away the surface features of our everyday practice and (2) highlight something in our practice which supports GJR. A familiar package of views attempts to satisfy both requirements. First, it is suggested that the structure of our usual justificatory activities can be explained in terms of epistemologically irrelevant practical and conversational features of the circumstances in which these activities take place. According to this explanation, we don't usually explicitly insist upon or attempt to satisfy GJR because doing so would be inconvenient and pointless, particularly since the justification would be more or less the same in every case and so can be taken for granted. Second, to support GJR this view offers the fact that GJR can seem compelling when we are engaged in epistemological reflection concerning our beliefs about the world. The idea here is that epistemological reflection removes the conversational and practical considerations which structure our ordinary justificatory activities, so that the purely epistemological requirements are

revealed. Consequently, if we insist upon GJR in this setting, then – it is claimed – GJR is correct.⁴⁶

This is an awkward package of views, if one also accepts minimal externalism about reasons. Consider, first, the explanation of the fact that GJR doesn't explicitly figure in our ordinary practice. A minimal externalist about reasons can explain this fact by suggesting that the structure of our justificatory practice reflects the fact that (because minimal externalism is true) our reasons for our beliefs about the world can only be local. This explanation is incompatible with the practical/conversational explanation of our ordinary practice, since in order to connect GJR with our practice the practical/conversational explanation must allow that there are certain circumstances in which the activity of justifying takes a non-local form or involves non-local requirements. Consequently, one faces yet more explanatory work if one accepts *both* GJR and minimal externalism about reasons: one has to explain why the truth of minimal externalism about reasons does not provide the correct explanation of the fact that our everyday justificatory practices involve only local justifications. It is not at all obvious how such an explanation might go. For instance, it might be suggested that our practice's structure is not explained by minimal externalism about reasons because we (implicitly) accept a different, incorrect, theory of reasons. But any such suggestion will face two further challenges. First, if minimal externalism is true, then what plausible explanation could be given for the putative fact that we (implicitly) accept an incorrect theory of

reasons? Second, what evidence is there that our practice is actually guided by some theory of reasons other than minimal externalism about reasons, given that nothing in our everyday practice appears to support GJR?

This last question brings us to the support for GJR which is alleged to come from the requirements imposed during epistemological reflection. Two things should be said about this. First, it simply is not true that we endorse GJR without reservation during epistemological reflection. Second, there is a plausible alternative explanation of our attraction, such as it is, to GJR. Recall the example regarding one's beliefs about physics discussed at the beginning of this section. As that example shows, our practice includes legitimate requests for general justifications regarding certain domains. This fact can explain the naturalness of the demand for a fully general justification of our beliefs about the world: we are tempted to generalize a form of justificatory request which is clearly legitimate regarding certain limited domains. However, if the requirements which govern our justificatory practice reflect the truth of minimal externalism about reasons, then it would be a mistake to assume that justificatory requirements of a sort applicable in the one case would also be applicable in the other. Minimal externalism about reasons can thus charge that the demand for a general justification of one's beliefs about the world is an inappropriate verbal generalization of a justificatory demand which is legitimate only regarding more specific domains. And this charge is buttressed by our ambivalence, even during epistemological reflection, about GJR.

I conclude that the advocate of minimal externalism about reasons can have good reasons for rejecting GJR. Accordingly, he or she can reasonably hold that our inability to satisfy GJR has no significance for the justificatory status of our beliefs; the reasons requirement can be met regarding each particular belief in virtue of one's ability to provide good reasons for holding it, where provision of the relevant reasons may well presuppose other background beliefs about the world. Accordingly, if we accept minimal externalism about reasons we can reasonably reject the general justificatory project, even if we accept the reasons requirement.⁴⁷

IV. Conclusion

Our ordinary activities of evaluating knowledge claims and justifying beliefs about the world always take place within a background framework of considerations about the world. If minimal externalism, and particularly minimal externalism about reasons, is true, then we can understand why this would be the case. We can also see why it should be philosophically satisfying. For I have argued that if an appropriate form of minimal externalism is correct, then there is no good intellectual motivation for undertaking a general vindication of our knowledge of the world or a general justification of our beliefs about the world. Consequently, if an appropriate form of minimal externalism is true, then – so far as I can see – we have no good reason to engage in an intellectual enterprise structured by the Cartesian Constraint.

One doubt may remain. Does the desire to provide a non-question-begging response to the skeptic give us a compelling motivation for the traditional project? This suggestion places a great deal of weight upon the argument in favor of skepticism, and though I cannot do full justice to the issue here, it is doubtful that there is a skeptical argument which can do the necessary work. If the Cartesian Constraint is supposed to arise from a concern not to beg the skeptic's question, then whatever argument is used to support skepticism must not itself presuppose that one may not make use of considerations about the world. Otherwise, the vantage point of the traditional project will have been in place from the beginning. This consideration places significant constraints on the kind of skeptical argumentation which could spur acceptance of the Cartesian Constraint. For instance, the argument cannot simply assume that our ultimate evidence for any given belief about the world cannot include other beliefs about the world. Likewise, the argument can't proceed simply by utilizing a requirement such as "In order to know p , you must know that you are not dreaming (or not a brain in a vat)." For unless one has *already* been forbidden from appealing to considerations about the world, one can appeal to considerations about the world as part of an argument that one meets this requirement. It is possible that there is a principled way to preclude such responses without simply assuming the vantage point of the traditional project. However, the antecedent implausibility of skepticism strongly suggests that whatever skeptical

argument does the trick will have other failings and can be countered without taking up the traditional project.⁴⁸

I conclude, then, that even if justification requires one to be able to provide reasons for one's beliefs, minimal externalism can enable us to reasonably refuse to continue to think in ways that are structured by the traditional epistemological project.⁴⁹

Bibliography

- Alston, William (1989). *Epistemic Justification: essays in the theory of knowledge*, Ithaca, NY: Cornell University Press.
- Annis, David B. (1978). "A Contextualist Theory of Epistemic Justification," *American Philosophical Quarterly*, reprinted in Paul K. Moser, ed., *Empirical Knowledge*, Totowa, NJ: Rowman and Littlefield, 1986.
- Audi, Robert (1993). *The Structure of Justification*, New York, NY: Cambridge University Press.
- Bealer, George (1999). "A Theory of the A Priori," *Philosophical Perspectives* 13, pp. 29 – 55.
- BonJour, Laurence (2002). *Epistemology*, Rowman & Littlefield.
- (1999) "Foundationalism and the External World," *Phil. Perspectives*, 13, 1999, pp. 229 – 249.
- (1998). *In Defense of Pure Reason*, Cambridge University Press.
- (1985). *The Structure of Empirical Knowledge*, Cambridge, MA: Harvard Univ. Press, 1985.
- Burge, Tyler (1993). "Content Preservation," *Philosophical Review*, 102, pp. 457 - 88.
- Cavell, Stanley (1979). *The Claim of Reason*, Oxford: Oxford University Press.

Dretske, Fred (1981). "The Pragmatic Dimension of Knowledge," *Philosophical Studies* 40, pp. 363-78.

Fumerton, Richard (1995). *Metaepistemology and Skepticism*, Rowman & Littlefield.

Gettier, Edmund. "Is Justified True belief Knowledge?" *Analysis* 23 (1963), pp. 121-3.

Goldman, Alvin (1999a). "Internalism Exposed," *Journal of Philosophy* 96, 271-93.

(1999b). "A Priori Warrant and Naturalistic Epistemology," *Philosophical Perspectives* 13, pp. 1 – 28.

Hare, R. M. "Philosophical Discoveries," *Mind*, LXIX, 1960, pp. 145 – 62.

Leite, Adam (Unpublished manuscript). "On Justifying and Being Justified".

(Unpublished doctoral dissertation). "Reasonable Doubts: Skepticism and the Structure of Epistemic Justification," Harvard University, 2000.

Plantinga, Alvin (1993). *Warrant: the Current Debate*, Oxford University Press.

Pollock, John and Cruz, Joseph (1999). *Contemporary Theories of Knowledge*, Rowman and Littlefield.

Pryor, Jim (2001). "Is Moore's Argument an Example of Transmission Failure?", unpublished manuscript, available at:

<http://www.princeton.edu/~jimpryor/papers/index.html>

(2001). "Highlights of Recent Epistemology," *Brit. J. Phil. Sci.* 52 (2001), 95-124.

(2000). "The Skeptic and The Dogmatist," *Nous* 34(4), 517-549.

Quine, W.V.O. (1986). "Reply to Stroud," *Midwest Studies*, pp. 473 - 455.

Sosa, Ernest (1994). "Philosophical Scepticism and Epistemic Circularity," *Proc. Arist.*

Soc. suppl. vol. 24, 263 - 90.

Stroud, Barry (2000). *Understanding Human Knowledge*, Oxford University Press.

(1986). "The Significance of Naturalized Epistemology," *Midwest Studies*, pp.

455 - 71.

(1984). *The Significance of Philosophical Scepticism*, New York, NY: Oxford

University Press.

Van Cleve, James. (1979). "Foundationalism, Epistemic Principles, and the Cartesian

Circle," *The Philosophical Review*, 88:1, 55 – 91.

Williams, Michael (1996). *Unnatural Doubts: Epistemological Realism and the Basis of*

Scepticism, Princeton, NJ: Princeton University Press.

Williamson, Timothy (2000). *Knowledge and Its Limits*, Oxford University Press.

¹ This characterization is intended to include the work of philosophers as diverse as Descartes, Locke, C. I. Lewis, and Laurence BonJour. For simplicity, I characterize the project in terms of an attempt to account for our *knowledge* of the world. However, some versions of the traditional project assume that knowledge requires justified belief and are motivated primarily by considerations about justification, while some are couched in terms of justification alone. For this reason, section III of my paper will focus on issues relating to justification. In the earlier sections my discussion will focus mainly on knowledge; footnotes will highlight parallel points regarding justification.

² Michael Williams has recently defended this answer as well (see particularly his 1996). For reasons which will be apparent below, however, I disagree with his view that the traditional project requires a problematic theory of justification (experiential foundationalism).

³ Bonjour (1985), pp. 56-7.

⁴ Bonjour (2002), pp. 236-7.

⁵ Stroud (2000, essays 8 and 10) and Fumerton (1995)

⁶ A few clarifications may be helpful. First, I make no commitments about the precise nature of the modality involved in this thesis, and nothing in my argument will turn on this issue. Second, minimal externalism about justification and minimal externalism about knowledge are distinct theses. I treat them together here merely for convenience.

(How one understands their relationship will depend in part upon how one understands the relation between knowledge and justification, an issue on which I will remain neutral for the purposes of this discussion.) Third, though it is now somewhat standard to understand the internalism/externalism issue in terms of epistemic accessibility requirements (Pryor, 2001), I focus instead on minimal externalism because it bears the most interesting and important relation to the traditional project. (I discuss the relation between minimal externalism and accessibility requirements below.) Finally, I don't want to argue over the phrase "the world outside of one's mind." The sorts of factors I have in mind are clear enough from paradigm "externalist" theories, such as basic reliabilist views of knowledge, Dretske's Relevant Alternatives theory (Dretske, 1981), and Alston's version of foundationalism (Alston, 1989, essay 9). So far as I can see, considerations relating to "externalist" approaches to thought-content such as those pioneered by Burge and Putnam have no direct bearing on the issues under discussion here.

⁷ Williams (1996), Annis (1978).

⁸ See, for example, Bonjour (2002), pp. 267-8.

⁹ I use the term "localist" rather than the more familiar "contextualist" in order to avoid confusion with current contextualist accounts of the semantics of "knows" and its

cognates. The view which I term localism need not be committed to a contextualist semantics for "knows" or any other term of epistemic appraisal.

¹⁰ According to the standard characterization, Gettier showed that justified true belief is not sufficient for knowledge. However, this characterization depends upon a certain interpretation or account of what justification is. I prefer to gloss Gettier's lesson this way: if you specify everything there is to be specified about what is going on inside the person's mind and the truth-value of the belief, you still haven't said enough to make it determinate whether the person has knowledge; you also have to say more about the world. That is just the thesis of minimal externalism about knowledge.

¹¹ Or, more weakly, in order to know or justifiably believe something, one must be in a position to determine through introspection and *a priori* reflection alone, of the conditions which suffice for knowledge or justified belief, that they obtain.

Minimal externalism and accessibility internalism (on either formulation) are independent theses. Since it is obvious that one could accept one and deny the other (or deny both), what needs to be shown is that one could accept both. If one accepts both, then (on the stronger formulation of accessibility internalism) one will hold that in order to know something, one must be in a position to know through introspection and *a priori* reasoning that the relevant conditions in the world are met. If one can gain knowledge of the world through introspection and *a priori* reasoning, then one *is* in such a position. So

the combination of minimal externalism and accessibility internalism is perfectly intelligible, if one also holds that knowing things about the world requires one to be able to gain knowledge of the world through introspection and a priori reasoning. This latter thesis is, of course, simply a version of traditional foundationalism.

¹² Three lines of thought are common in the literature. The first attempts to derive the accessibility requirement from the requirements of epistemic responsibility, where responsible belief is taken to require that one have (and recognize) some reason for thinking that one's belief is true (see, for instance, Bonjour (1985), chapter 1). But epistemic responsibility, so understood, does not entail the accessibility requirement (on either formulation); the requirement that one recognize a reason for thinking one's belief true does not entail *anything* about the ability to find out about whether one is justified or has knowledge (or even to find out, of the conditions which suffice for knowledge or justified belief, that they obtain). The second line of thought attempts to derive the accessibility requirement from the idea that it is constitutive of the concept of justification that considerations about justification can play a regulative role in deliberations about what to believe. Again, however, this idea does not successfully motivate either version of the accessibility requirement; one could accept it and hold that judgments about justification depend upon considerations about the world which could not be arrived at through introspection and a priori reflection alone (see Goldman,

“Internalism Exposed” (1999a)). Indeed, one could come to accept such a position by combining certain forms of minimal externalism with a regulative conception of justification. (If one is already thinking about matters within the terms of the traditional project, then accepting the regulative role of justification will appear to force acceptance of the accessibility requirement. But then considerations about the regulative role of justification won't be what motivates the Cartesian Constraint (via the accessibility requirement); the Cartesian Constraint will have been in place from the beginning.) Finally, a third line of thought attempts to derive the accessibility requirement from the idea that being justified is a matter of having fulfilled one's epistemic duties (a so-called "deontological" conception of justification) (see, for instance, Plantinga (1993), chapter 1). Since one cannot be held responsible for meeting (or failing to meet) one's duties unless one is able to find out what they are, the justificatory status of one's beliefs must be something one can find out about. Again, however, one could perfectly well accept this line of thought and reject the accessibility requirement (on either formulation). One might be able to find out about one's duties, but not by means of introspection and a priori reflection alone. This is often the situation with regard to moral duties.

¹³ See, for instance, Ernest Sosa, (1994).

The accessibility requirement and the reasons requirement are completely independent. The former does not entail that knowledge or justification requires one to

have adequate reasons for one's belief. And the reasons requirement allows that one can be justified (or have knowledge) and yet not be able to find out, through introspection and a priori reflection alone, about some of the factors in virtue of which one is justified (or has knowledge).

¹⁴ One other form of "internalist" requirement may seem relevant here. It is sometimes suggested that the traditional project is motivated by the so-called "KK principle" (If you know that p, then you know that you know it) or by its correlate for justified belief.

However, both principles are manifestly false. Moreover, neither principle actually does motivate the Cartesian constraint. For simplicity, consider just the KK principle (the same point applies to its correlate for justified belief). This principle leaves open *how* one knows that one knows the proposition in question. So even if one were inclined to accept some such principle, one could hold that one has the relevant second-order knowledge in part because of things which one knows about the world. The principle consequently would not generate the Cartesian constraint, even if it were not otherwise defective.

¹⁵ Stroud, 2000, p. 120.

¹⁶ This requirement allows the details of the explanation to differ in different cases, but it rules out accounts which ultimately amount to a disjunction of two or more completely disparate explanations.

¹⁷ This project may be more or less metaphysically loaded, depending upon whether some kind of reductionist or supervenience thesis is presupposed.

¹⁸ Ernest Sosa correctly points this out in the course of arguing that externalist views can provide a satisfying general explanation of our knowledge of the world (1994, p. 288).

¹⁹ 1969 Notre Dame lectures, in *Essays in Radical Empiricism*, p. 277. Firth sharply distinguishes his project from any attempt to justify our empirical beliefs in general. Developing the reconstruction does not provide justification for our beliefs, but rather exhibits relations and principles of inference in virtue of which our justified beliefs are already justified (p. 277).

²⁰ Consider, for example, the traditional experiential foundationalist form of reconstruction favored by Firth and Chisholm. Firth suggests that the transition from beliefs about one's experience to beliefs about objects in an external world could be justified by an a priori justificatory principle "that anyone who has a particular kind of sense-experience has '*prima facie* evidence' that he is perceiving a physical object that has the correlated objective property" (pp. 311-12). This principle could be replaced with the following: "Given how the world in fact works, and given the nature of our sensory systems, having a particular kind of sense-experience is a good *prima facie* indication that there is a physical object having the correlated objective property." (This is, more or less, William Alston's view (1989, chapter 9).) On such a view, our possession of

justified beliefs about the world would depend upon how things in fact tend to go in the world, and the general reconstructive explanation of our knowledge would simply reflect this fact by making use of claims about the world.

²¹ Stroud, "The Significance of Naturalized Epistemology", pp. 80-1.

²² Concerning Quine's proposal to view our beliefs about the world as a construction or projection from sensory stimulations, Stroud writes, "If we tried to think of *all* of our own beliefs as a 'construction or projection from stimulations', we would at most have access to what we know to be our assertions or beliefs about the world, but we would not in addition have independent access to the world they are about on the basis of which we could determine whether they are true" (op cit.). I am not sure what Stroud means by "independent access to the world." Regardless, however, it is not true that if we understand our own beliefs in the suggested way, it becomes illegitimate to make explanatory use of what we know of the world. One can simultaneously understand one's beliefs as arising from a process of "construction or projection" *and* make use of what one knows about the world in order to explain how one has gained knowledge – so long, that is, as one has not suspended one's commitment to the truth of those beliefs.

²³ Though I focus here on a question regarding knowledge, a version of the traditional epistemological project can be generated from a parallel question about justification.

Laurence Bonjour has recently appealed to this sort of reflective curiosity to motivate his version of the traditional epistemological project: “One ... question that can be asked right now is whether I really do know all of the things that I think I do ... [or] whether I really know any of them at all” (2002, p. 5). He concedes that there may not be any particular reason to suspect that we *don't* know any of the things we take ourselves to know, but urges that considering the question is “intellectually valuable” (p. 6). Richard Fumerton depends upon similar considerations in arguing that externalist approaches are not philosophically satisfying (1995, p. 177). Stanley Cavell, too, stresses the role of something like this question in generating the intellectual demands which lead to the traditional epistemological project (1979, pp. 140ff.).

²⁴ For a provocative recent discussion of the relation between assertion and knowledge claims, see Williamson (2000, chapter 11).

²⁵ Moreover, contrary to Michael Williams’ claim (1996) the project doesn’t require foundationalist assumptions in order to motivate grouping together our beliefs about the world as a class worthy of epistemological investigation. These beliefs can be classed together simply in virtue of what they are about: contingent matters involving objects, events, etc., which do not depend upon anyone’s mental states for their existence or for being as they are (Stroud, 2000, essay 9).

²⁶ Stroud, "Understanding Human Knowledge in General" and "Scepticism, 'Externalism,' and the Goal of Epistemology," both in his (2000).

²⁷ Sosa, for instance, offers a version of this response:

[I]t is ... unclear why the epistemologist needs to *see himself as having* justification for his theory, or as knowing his theory, in order for it to give him understanding of how he and others know the things they know, either in general or in the domain in question. Why is it not enough that he in fact *have good reason to accept his theory* or perhaps even *know his theory to be true*? This is different from his knowing that he has good reason to believe his epistemologically explanatory theory, or even knowing that he knows his theory to be true (1994, p. 275, italics in original).

²⁸ As mentioned in fn. 23 above, a version of the traditional epistemological project can be motivated by raising a general vindicatory question regarding the justificatory status of our beliefs about the world. Though my discussion in this section will focus on knowledge, the argument equally applies, *mutatis mutandis*, to this general vindicatory question about justification (provided it is not assumed that whether one's beliefs about the world are justified depends upon one's ability to provide a general justification for them).

²⁹ It might seem that certain forms of minimal externalism would avoid this consequence.

Suppose that one could make a reasonable, justified, or knowledge-giving inference from claims about the nature of one's sensory experiences, apparent memories, and current beliefs, as well as any relevant deliverances of a priori intuition (if there is such), to some claims about the world. Couldn't a minimal externalist then appeal to those claims about the world as part of a successful general vindicatory explanation of her knowledge of the world? No. The resulting theory would not have the desired generality. To see this, consider the claims about the world which are supposed to be inferred from the more restricted basis. Call them A. One's knowledge of A is part of what the theory is supposed to explain. Suppose, then, that a minimal externalist attempts to give a general explanation of his knowledge of A. It will look like this:

(1) I know that A because I inferred A from [the more restricted basis], and

(2) Conditions CW are satisfied with regard to my beliefs about A.

(2), however, is a claim about the world, so the minimal externalist's knowledge of its truth is part of what the theory is supposed to explain. And given his commitment to minimal externalism, *any* attempted explanation will also involve such a clause appealing to circumstances in the world. An appeal to circumstances in the world is thus an unavoidable feature of the minimal externalist's attempt to pursue the general vindicatory project. Consequently, a minimal externalist simply can't do what the general vindicatory

project requires. (This point does *not* provide an argument against foundationalism or any other theory of justification. It merely shows that if one accepts minimal externalism, then one will not be able to answer the general vindicatory question *regardless* of what account of justification one accepts.)

My argument here is indebted to a line of thought suggested by Stroud's "Understanding Human Knowledge in General" and "Scepticism, 'Externalism,' and the Goal of Epistemology," both in his (2000). Stroud seems to take this line of thought to be a mark against externalism, whereas I will argue it provides the minimal externalist with good reasons for rejecting the general vindicatory project. (Stroud also fails to distinguish the general vindicatory project from the other general explanatory projects.)

³⁰ This point assumes that having knowledge about the world does not require being able to satisfactorily complete the general vindicatory project. This assumption seems quite reasonable. Moreover, since the general vindicatory project does not itself depend upon such a requirement, we can unproblematically rely upon this assumption at this point in the argument. I am grateful to Fred Schmitt for raising this issue.

³¹ There may be emotional or "existential" motivations for taking the general vindicatory question seriously. Minimal externalism does not address this issue, and I will not take it up here. Such issues are addressed in interesting ways in the work of Stanley Cavell.

³² Laurence Bonjour puts the issue this way:

My ultimate concern is with what reasons there are for thinking that our familiar beliefs about the physical world are true, where I have in mind reasons: (1) that do not in some way beg the question by presupposing the acceptability of other beliefs about physical objects; and (2) that are at least in principle available through reflection and analysis to believers more or less like ourselves (1999, at p. 229).

³³ If one denies that knowledge requires justification, then one can reject the general justificatory project from the start, at least insofar as it is supposed to have any relevance to questions about knowledge. However, since the traditional project can be framed in terms of justification just as well as in terms of knowledge, the relation between minimal externalism and the general justificatory project would still remain of interest.

³⁴ A fully defensible version of this principle would have to be hedged and elaborated in various ways. For details, see my unpublished manuscript, “On Justifying and Being Justified.”

³⁵ See, for instance, Pryor (2001), Audi (1993, p. 145), Alston (1989, p. 82, p. 83 fn. 3), Sosa (1994, where he explicitly links his view with the aim of turning aside the global justificatory project.)

³⁶ I defend the reasons requirement in my manuscript, “On Justifying and Being Justified”.

³⁷ Minimal externalism about justification would be a wider view than what I'm calling "minimal externalism about reasons," since it would include process reliabilist views as well as views which link justificatory status with the adequacy of the believer's reasons. However, like minimal externalism about justification, minimal externalism about reasons does not require the reasons requirement, since one could hold that one can *have* good reasons without *being in a position to provide them*. (One might think, for instance, that very young children and perhaps some animals are in precisely this situation.) Alston holds such a view, though he generally speaks in terms of "grounds" rather than "reasons" (1989).

³⁸ Alston (1989, essay 9). Again, since Alston denies the reasons requirement, he would reject the demand for a general justification from the get-go. I appeal to his view here only to provide an example of a view which accepts minimal externalism about reasons.

³⁹ This claim is weaker than the following disputed thesis:

In order to successfully justify one's belief that *Q* by appealing to *P*, one must have some prior and independent reasons in favor of the proposition *that P is a good reason for believing that Q*.

Although I am strongly tempted by this thesis, nothing in my argument here requires it.

⁴⁰ I take this last point to be a truism. However, it can appear to be in tension with an approach to justification which has been favored by some recent defenders of

foundationalism. According to this approach, the following situation is possible: (1) P's obtaining is a good reason for believing Q, (2) one is consequently justified in believing that Q in virtue of the fact that P, and (3) (in part) in virtue of being justified in believing that Q, one is justified in believing that P is a good reason for believing Q. (This possibility is explicitly allowed for by Pryor (2001, unpublished manuscript) and Van Cleve (1979). Alston's view in his (1989) allows for it as well, and he explicitly makes use of it in his (1986).) It would seem that in such a situation one could legitimately attempt to justify one's belief that P is a good reason for believing Q by arguing as follows: "P; so Q; ... ; so P is a good reason for believing Q." Consequently, this view seems to permit a form of justificatory argument which I've claimed is clearly unacceptable. However, an advocate of this view need not disagree with my claim, because an advocate of this view can divorce *being justified* from *being able to provide a successful justification*. Pryor and Alston both do so, as noted above. Accordingly, they could accept my claims about the structure of the activity of justifying, since these claims carry implications about justificatory status only if one *also* accepts, as they do not, that to be justified, one must be able to provide a successful justification. (The consequence of their view, which may be troubling, is that one could *be justified* in virtue of considerations to which one *could not* acceptably appeal as part of a justifying argument.)

⁴¹ Note the conditional nature of this claim. I have argued that *if one accepts an appropriate form of minimal externalism*, then one cannot provide a general justification for one's beliefs about the world. I haven't made any claims about whether the project can succeed if one rejects minimal externalism. Consequently, it may be structurally possible to provide a general justification for one's beliefs about the world if one holds — as Chisholm, for instance, does — that the relevant linking propositions can be known or justified a priori. For Chisholm is not a minimal externalist about reasons; he holds that the a priori principles of evidence do not concern objective probability and are not otherwise determined by what is the case or "how things go" in the external world. (For more on the relation between minimal externalism about reasons and the claim that relations of evidential support are a priori, see fn. 43.)

⁴² Quine and many other contemporary naturalists accept at least this view.

⁴³ What if one accepts minimal externalism but also holds that the relevant support relations can be known or justified *a priori*? First, it should be noted that it is possible to be a minimal externalist regarding the *a priori*. According to one important characterization, one has an a priori justification just if one has an appropriate purely intellectual "seeming": it strikes one, in an intellectual way, that such-and-such is the case. (For two ways of developing this idea, see Bealer (1999) and Bonjour (1998).) A minimal externalist about reasons could hold that it is in virtue of facts about the world

that a purely intellectual seeming is in fact a good indication of the truth of a certain proposition. (Alvin Goldman considers a parallel position within a reliabilist conception of a priori justification in his (1999b).) If one held such a view, then one could not acceptably appeal to intellectual seemings as reasons for holding beliefs about evidential support relations unless appropriate background propositions about the world were not in the suspension set. Consequently, one could not provide a fully general justification of one's beliefs about the world. Any other version of minimal externalism about a priori justification will have the same result.

There may be a version of minimal externalism about reasons which, if it were otherwise successful, could enable one to provide a general justification for one's beliefs about the world. This view would (a) deny minimal externalism about a priori knowledge and justification, (b) hold that we can have a priori knowledge or justification regarding propositions about the goodness of reasons for believing claims about the world, and yet (c) maintain (with the minimal externalist) that the goodness of such reasons is determined by what is the case and how things tend to go in the world. For all I've said here, this view might succeed in providing a general justification, since it would allow that even in the context of the general justificatory project one could acceptably appeal to principles of evidential support to the effect that if certain things are true internal to one's mind, then it is highly likely (in some appropriately objective sense) that

certain propositions about the external world are true. I cannot fully evaluate this sort of heroic rationalism here. However, even if it can be made coherent, the position strikes me as extremely implausible. How, exactly, are we supposed to gain this knowledge of how things go in the world? As Hume showed, just thinking really hard won't give us knowledge (or even justified beliefs) about causal connections or (more generally) about evidential connections that hold only if the world is a certain way. For this reason, it is tempting to follow Chisholm in thinking that one must choose between a priori principles of evidence and minimal externalism. (I am grateful here to Mark Kaplan.) For my purposes here, however, it is not necessary to refute this form of rationalistic minimal externalism about reasons, since it suffices that there are other, more plausible forms of minimal externalism which do have the upshot discussed in the main text.

As I noted earlier, Chisholm's move may make it structurally possible to engage in the general justificatory project. Nothing I have said here is intended to suggest otherwise.

⁴⁴ In particular, engaging in "analysis" of the concept *justified belief* seems fruitless. If concepts are understood as abstract objects akin to Platonic Forms, then the advocate of GJR could simply disagree with any opposing analysis, and I see no way to resolve the dispute unless we appeal to particular judgments instantiating our practice. Likewise, the advocate of GJR could always claim that an opponent has latched onto the wrong

concept; again, the only way to decide the dispute about which concept is “ours” is to look to our practice. And if we take concepts as psychological items, then I have mine and the advocate of GJR has his. I don’t see anything more that can usefully be said.

The only firm data, so far as I can see, are what we find in our practice.

⁴⁵ In such an investigation, we are like someone who tries to gain reflective understanding of how a certain dance is to be done by paying attention while doing the dance. (I allude here to R. M. Hare’s discussion of dancing the eightfold reel in “Philosophical Discoveries.” This paper was called to my attention by Gary Ebbs’ *Realism and Rule-Following*.)

⁴⁶ Stroud explores a package of views of this sort (concerning, in particular, the requirements for knowledge) in his 1984, chapter 2.

⁴⁷ Minimal externalism about reasons also motivates rejection of a closely related general question which has guided some traditional work in epistemology. The question is this: “Of all the things I believe about the world, which *should* I believe?”, where the question is supposed to be posed in such a way that an answer to it would not presuppose the truth of any claims about the world. (A version of this question is used to motivate the somewhat idiosyncratic version of the traditional project undertaken in Pollock and Cruz (1999).) The relevant sense of “should” is the epistemic one, where what one should believe is determined by the reasons for and against the truth of particular claims. An

argument parallel to the above enables us to reject this question as both unanswerable and epistemologically irrelevant. If minimal externalism about reasons is correct, then this question simply puts us in a position in which we can no longer make use of whatever good reasons we possess. But that position is irrelevant to the question of what we should believe about the world. If minimal externalism about reasons is correct, then reflection regarding what one should believe about the world is local.

Laurence Bonjour has indicated to me in personal conversation that he now believes the traditional project is best motivated simply by reflecting upon the question, “Are any of my beliefs about the world actually true?” Since the attempt to answer such a question would have to proceed in terms of the consideration of reasons, it seems to me that this question will collapse into the question, “Which, if any, of these things should I believe?”

⁴⁸ It might be suggested that by providing a scenario (such as the dreaming hypothesis) in which everything “inner” is just as it is and yet all of one’s beliefs about the world are false (or not knowledge), the skeptic somehow shifts the burden of proof onto us.

However, the mere fact that we can describe such a scenario has no tendency to show that we lack knowledge of the world, unless it is assumed that in order to have any knowledge of the world one must be able to construct an adequate argument, using only “inner” resources, in favor of some particular claims about the world. But that assumption is a

version of precisely the view that is at issue, so it seems that the standpoint of the traditional project appears here as a presupposition, not an upshot. For a detailed discussion of these issues, see Michael Williams (1996) and my unpublished doctoral dissertation. (Stroud has suggested more generally that the traditional project is a source, not a consequence, of external world skepticism (2000, essays 8-10, and 1984, especially chapter 7). Michael Williams (1996) takes a similar line, though he maintains, against Stroud, that the traditional project depends upon problematic philosophical presuppositions.)

⁴⁹ I would like to thank audiences at Indiana University, Bloomington, and the University of Wisconsin, Milwaukee, and, in particular, Laurence Bonjour, Mark Kaplan, Stephen Leeds, Fred Schmitt, Jonathan Weinberg, and an anonymous referee for *Philosophy and Phenomenological Research*.