

The Problem of the Criterion¹

Introduction:

The problem of the criterion in epistemology raises certain fundamental questions concerning the methods a philosopher ought to use in arriving at both analyses of epistemic concepts and conclusions concerning legitimate principles of reasoning and justification. The answers one gives to these questions are thought to bear directly on the challenge of skepticism. The outcome of battles is often determined by the ground occupied by opposing forces. Interest in the problem of the criterion is often a function of the desire to occupy favorable ground in the battle with the skeptic.

Despite the title of this paper, there are, I believe, at least two different problems one might associate with “the criterion.” One is, perhaps, best known through the work of Roderick Chisholm, whose various editions of *Theory of Knowledge* all devote a chapter to it. As Chisholm develops the problem, it concerns methodological questions concerning how to identify *sources* of knowledge or justified belief. But there is a closely related question concerning how to identify necessary and sufficient conditions for the application of epistemic *concepts* like knowledge and justified belief. In the following discussion I’ll distinguish these two methodological issues, and I’ll begin by looking at the question of how one should go about discovering the correct analysis of epistemic concepts.

Metaepistemology and the Problem of the Criterion:

Metaepistemological questions seek to uncover through philosophical analysis the nature of fundamental epistemic concepts or properties. In metaepistemology, we are interested in getting an account of what knowledge, evidence, justified belief, and so on are. To be sure, philosophers don’t agree with each other about just exactly what is involved in coming up with a correct philosophical analysis. Many analytic philosophers in the first part of the twentieth century thought that their task was to discover *analytically* necessary and sufficient conditions for the application of concepts. So if we were concerned with analyzing knowledge, our task was to discover some set of conditions each of which is analytically necessary for knowledge and all of which together are analytically sufficient for knowledge. The classic test for whether a condition for knowledge, say the truth of what is believed, is analytically necessary is whether or not it is absolutely *inconceivable* that someone has knowledge while failing to satisfy the condition. The test for whether a conjunction of conditions *X* is jointly sufficient for knowledge is whether we can conceive of *X* obtaining without knowledge.²

¹ I’d like to thank the faculty (particularly Larry BonJour and Ann Baker) and the graduate students at the University of Washington for their valuable criticisms and comments on an earlier draft of this paper. Michael Bergmann’s and John Grecco’s comments were also most helpful. Although I’m not even sure he will remember it, I’d also like to thank Richard Feldman for an interesting conversation we once had on some of the issues discussed here.

² I carefully talk of conceivability as a *test*. The modal operator implicitly contained in talk of what is *conceivable* precludes, I believe, an analysis of any sort of possibility in terms of conceivability.

Even philosophers who are more or less committed to the central role of thought experiments in testing proposed analyses, disagree with each other profoundly over just exactly what is the *object* of an analysis. So some thought that they were analyzing the meanings of words or sentences, some thought they were analyzing concepts or thoughts into their “parts,” and still others thought that they were analyzing properties, states of affairs, or facts, into their constituents.

These days matters are even more complicated given now fashionable externalist theories of reference and meaning. At least some externalists toy with the idea of modeling philosophical analysis on scientific “analysis.”³ The scientist investigates the world in an attempt to uncover the fundamental nature of various “natural” phenomena—water, heat, or lightning for example. Water turns out to be H₂O, heat turns out to be molecular motion, and lightning turns out to be an electrical discharge. No-one thinks that it was possible to discover these familiar facts about the natural world through thought experiments. Would that science were that easy. Still, many direct reference theorists will argue, the truths described above are in some important sense necessary. The molecular structure of water, for example, is not just a contingent feature of water. It is an essential feature of the stuff we call water that it has that molecular structure. Stuff that looks, tastes, and smells like water wouldn’t be water, the argument goes, unless it had the right molecular structure. If we can discover the essential nature of so-called natural kinds like water through empirical investigation, then perhaps we can discover the essential nature of knowledge, or justified belief, through empirical investigation.

Methodological Internalism and Analysis:

Given the radical differences between these conceptions of philosophical analysis, we must be careful in characterizing the problem of the criterion as it relates to the way in which one tests proposed analyses. Let’s begin, however, presupposing the more traditional conception of analysis as conceptual analysis—the sort of task a philosopher could complete from the arm chair, employing the method of thought experiment. Presupposing this framework, we are trying to discover necessary and sufficient conditions for the application of epistemic concepts. Let’s use the concept of knowledge as our example. Put starkly, the problem of the criterion is this. Does one first decide what one knows and then try to learn from paradigmatic examples of knowledge the necessary and sufficient conditions for knowledge? Or does one discover first the necessary and sufficient conditions for knowledge and apply what one learns to discover what one knows? Or perhaps, should one conclude that it is a mistake to suppose that either sort of discovery takes priority over the other? Following Chisholm’s terminology, we can describe those who give priority to particular knowledge claims, particularists. Those who insist that we be guided by antecedent conclusions about the correct analysis of knowledge, we can call generalists or methodists. The compromise position is most naturally associated with the method of reflective equilibrium made most famous by Rawls, but I’ll reserve discussion of this view until my concluding remarks.

The questions posed above are difficult to answer. Each approach may seem *prima facie* plausible. On the one hand, if we look at how philosophers actually do test proposed analyses of knowledge (or other epistemic concepts), they do seem to rely

³ See Kornblith (2002)

critically on the method of counterexample and “intuitions” about particular cases. There is no argument form used more frequently by philosophers than *modus tollens*, and the kind of *modus tollens* most favored by philosophers is the *reductio ad absurdum*. A *reductio* is just *modus tollens* with appropriate histrionics employed in the denial of the consequent—If *P* then *Q*, but (for God’s sake anyone in his right mind knows that) not-*Q*. So not-*P*. And nowhere is the *reductio* employed more often than in the attempt to argue against proposed philosophical analyses. In *The Republic* (331c) after Socrates takes Cephalus to be suggesting that right conduct is just telling the truth and returning what you owe to others, Socrates immediately rejects the view by pointing out that you obviously shouldn’t return weapons you’ve borrowed from a man who has gone mad. All kinds of twentieth century epistemologists thought that we have knowledge whenever and only when we have justified true belief. Gettier (1963) described a couple of cases in which people have justified true beliefs but don’t have knowledge and most everyone abandoned the view.⁴ So it certainly looks as if our search for correct analyses of epistemic concepts is guided by particular judgments we make about what does and doesn’t satisfy the concepts in question.

Still, one must wonder how philosophers who lack the correct analysis of knowledge are in a position to decide what people do or don’t know. Furthermore, on the traditional view one is supposed to be able to know *a priori* the necessary and sufficient conditions for knowledge. That we know this or that proposition is always a contingent matter. But *a priori* knowledge is usually thought to be more secure than knowledge of contingent truths. So why would the traditional analytic philosopher let *a priori* knowledge of analytic truths be held hostage to contingent truths about knowledge of particular cases?

Again, one might suppose that the position one takes on this issue has significant implications for the way in which one approaches the challenge of skepticism. If we emulate Moore and assume at the outset that we know all sorts of particular truths about the past, the physical world, and the future, using that assumption to guide our analysis of knowledge, the skeptic obviously can’t win. We *start* with the rejection of skepticism, and we defend that starting point by arguing that philosophical methodology requires it. Philosophy can’t proceed without a healthy sense of what is absurd, and radical skepticism, the argument goes, is clearly absurd.

The details of our defense of the particularist starting point will be shaped by our particular metaphilosophical view of analysis. For example, if we think that analysis is a matter of discovering the meanings of words, we might defend the particularist’s starting point by arguing for the slogan that meaning is use and by pointing out that we do in fact use the term “know” frequently and unapologetically to describe our relation to all sorts of truths. To offer an account of knowledge that makes all or most paradigmatic knowledge claims false is to exhibit a failure to understand what it is for an expression to have meaning. Paradigmatic uses of an expression *determine* its meaning.

However attractive the particularist approach to traditional philosophical analysis might initially seem, at least one version of the view doesn’t stand up to close examination. It is, in fact, a bit misleading to suggest that on the traditional view one evaluates one’s proposed analyses of epistemic concepts by testing them against

⁴ Though not, of course, the spirit of the view. There continues the (seemingly never-ending) search for a Gettier-proofing condition to add to the traditional analysis of knowledge.

“intuitions” about which particular propositions we know or are justified in believing. We have already had occasion to note that conceptual analysis is guided by *thought* experiments. Gettier didn’t wander around looking for someone who believed falsely that Jones owned a Ford while randomly forming a belief in a disjunctive proposition involving Brown and Barcelona. He asked us to *imagine* a situation and filled in the details of that imaginary situation as he pleased. The nice thing about thought experiments is that one gets to assign the properties of the imagined situation as one will. Now to be sure, we do seem to discover *something* in thinking about the hypothetical situation. We discover, I would suggest, something about how we use the term “know”—we discover something about our linguistic dispositions. Specifically, we discover that we wouldn’t describe the particular situations Gettier describes as cases of knowledge.

I have argued elsewhere (1983) that it is a mistake to suppose that one can divorce philosophical analysis from language. The idea that when one is engaged in philosophical analysis one begins by getting before one’s mind a non-linguistic concept, property, state of affairs, or fact, and then performs a philosophical dissection of it, is phenomenologically wildly implausible. If that’s how philosophical analysis were to proceed, it would be utterly mysterious as to how philosophers come to such radically different positions concerning the analysis of knowledge, causation, justified belief, moral obligation, and so on. Rather, in trying to discover what knowledge, say, is, I think the philosopher is trying to discover facts about which possible situations would be correctly described using the term “know.” More precisely still, I think that each philosopher is trying in the first instance to discover facts about his or her own linguistic dispositions—conclusions that we typically presume will inform others as to *their* linguistic dispositions.

I can’t convince you here of my own view of philosophical analysis. The first point I want to stress, though, is that however one thinks of *traditional* philosophical analysis, one won’t, one shouldn’t, test analyses against particular *empirical* claims about what is known or not known. To convince yourself of this, suppose that you are actually in a Gettier situation. You are looking at the broken clock that reads 3:00 p.m. just as it happens to be 3:00 p.m.⁵ Because you don’t know that the clock is broken, you naturally enough form the true belief that it is 3:00. Now also since you don’t know that you’re in a Gettier situation, I suppose that if you were inclined to make Moore-style pronouncements, you might insist that this is a clear example of knowledge and that it better turn out that any plausible analysis of knowledge implies that you know this proposition. Clearly, though, were you to discover relevant contingent features of your actual situation, you would recant your insistence that this turn out to be a case of knowledge. The moral to draw is that you never really were relying critically on *this* being a case of knowledge. The most you were doing in using this example was to conjure up a *kind* of hypothetical situation of which you took this particular example to be an instance. But it didn’t matter whether it was or wasn’t an instance of the kind. It was the hypothetical situation—the situation *thought of* in a certain way—against which you were testing your intuitions. And the *conditional* statement, if a belief were to satisfy the description in question it would be knowledge, if true, is every bit as analytic as the general statement, *S* knows that *P* iff *X*.

⁵ Russell’s (1948, p. 154) example given long before Gettier.

Indeed, once it becomes apparent that it is *subjunctives* the truth value of which is guiding our proposed analysis, it is not at all clear that we can still so easily draw the distinction between the generalist who insists on the priority of knowledge of philosophical analyses, and the particularist who insists on the priority of knowledge of particular truths about what we know. Even if subjunctive conditionals contain expressions that appear to pick out particular objects, it is often the case that their truth conditions are just as general as the truth conditions of general propositions. Suppose, for example, I tell you that if this (pointing to my pencil) were metal, then it would expand when heated. In one sense, it looks as if the claim is one about a particular object—the pencil. But further reflection surely suggests that it is more plausibly construed as a claim about lawful or causal connection between properties. After all, the fact that I picked out the pencil in making the claim is utterly irrelevant to the truth of the subjunctive. I could have pointed to my phone, my arm, the water in my glass, or anything else for that matter, and it wouldn't have affected in the slightest the truth of my statement. What makes it true that if the pencil (my phone, my arm, the water in my class) were metal it would expand when heated, is the general truth that metal expands when heated. I'm suggesting that the truths involving hypothetical situations that guide our claims about how to analyze knowledge are like the above "particular" subjunctive truth about the pencil. If it is true that a particular belief would count as knowledge if and only if it possesses characteristics *X*, what makes it true will be the fact that we can analyze knowledge as beliefs with characteristics *X*.

Still, the above account seems to leave mysterious the force that counterexamples have for philosophers. It certainly seems that we often do find a given analysis of some concept plausible, only to be convinced by consideration of some *more* detailed hypothetical situation that our initial view was mistaken. Remember the Gettier counterexamples. So what's going on? The best way to answer the question, it seems to me, is to think carefully about what happens when you do get convinced by a counterexample. You start out, let's say, thinking that you would be prepared to describe as knowledge all and only strongly justified true beliefs. You then think of a possible situation in which it is pretty obvious that a strongly justified true belief is present, even though you wouldn't describe it as a case of knowledge. What you discover, it seems to me, is that you don't have the disposition that you thought you had. That discovery of course, doesn't tell you what features of the situation entertained are prompting your refusal to characterize the belief in question as knowledge. There are all sorts of properties posited in the description of the situation. You might have withheld the judgment that the belief is knowledge because the situation entertained was one in which you reasoned through falsehood. Or you might have been influenced by the fact that the imaginary situation involved defeaters. Or perhaps the situation might have reminded you forcefully that your justification didn't satisfy appropriately strong Cartesian requirements for knowledge. And you needn't be *consciously* aware of any of these features of the posited situation in order for them to play a causal role in your withholding an ascription of knowledge. That you had only implicitly filled out the imaginary situation with certain properties present seems to leave unaffected the capacity of their presence to affect your judgment. Just as we can feel uneasy about some action we are thinking about undertaking without understanding what makes us uneasy, so also, it seems that we can realize that we are uneasy about calling a given belief knowledge even if we aren't sure

what is making us uneasy. Again, when consideration of some hypothetical situation affects us in this way, we come to realize that we don't follow the semantic rule we thought we followed governing knowledge.⁶

So on the traditional conception of analysis, it seems that we can discover facts about our dispositions that are relevant to the truth of our proposed analyses by careful consideration of thought experiments. We can, if we like, call this a version of particularism. It is important to stress, again, however, that it is a *thought* experiment that plays the evidential role in testing the analysis. It is *not* empirical conclusions concerning our actual knowledge or justified belief about some particular proposition that guides our analyses.

Methodological Externalism and Analysis:

One might suppose that the above arguments will leave untouched the new “particularists” who are determined to model philosophical analysis on scientific discovery. But considerations raised above tell against this conception of analysis as well. To pursue the scientific model, one surely must begin with “samples” of the epistemic kind we are going to empirically investigate. One can't perform an empirical assay of a kind without having samples of the kind to assay. So just as on one caricature of the view, we run across some relatively clear, odorless, liquid in bodies we call rivers and lakes, presume that it is all of the same kind, and proceed to investigate empirically the kind, so also, I suppose, the idea is that we run across some true beliefs that we take to be of the same fundamental kind and begin to empirically investigate the nature of that kind.

To state the view clearly is to see the obvious difficulties it encounters. How precisely are we to pick out the samples? The literature suggests at least three possibilities: 1) Our reference to knowledge, justified belief, evidence, and the like, might piggyback on the successful reference of others to such “stuff”—perhaps the experts in the field; 2) We might “fix the reference” of expressions such as “knowledge” and “justified belief” with definite descriptions, where the samples to be investigated are those things denoted by the definite descriptions; 3) We might demonstratively refer to the samples we are interested in investigating as we run across them. Consider each in turn.

We surely can't rely on experts to pick out the samples of epistemic properties the nature of which we want to empirically investigate. Who are the experts to whom we would defer when it comes to the referents of such terms as “knowledge” and “justified belief?” Would they include Descartes? Or perhaps, Hume? Or maybe Goldman? The problem, of course, is that these philosophers don't agree with each other on what should be included in the class of knowledge and justified belief. And no self-respecting epistemologist is going to defer to another philosopher when it comes to presuppositions about what counts as a paradigmatic instance of knowledge or justified belief.

⁶ Compare the thought experiments a linguist employs in testing hypotheses about the syntactic rules we follow in speaking a language. There is a marked similarity between these thought experiments and the thought experiments that test hypotheses about semantic rule following. Just as consideration of a given sentence can reveal that we don't follow a generalized syntactic rule, so also consideration of a given hypothetical situation can reveal that we don't follow a generalized semantic rule.

I have argued elsewhere (1989) that the very idea of reference fixing is an illusion. We can assign a name (for a particular or a kind) the meaning of a definite description, but we can't wave a definite description magically over a symbol and create thereby a "pure" referring expression. But suppose (contrary to fact) that we can use a definite description to mystically breathe referential life into an expression like "knowledge" or "justified belief" and that we can do it without assigning the terms the meaning of the definite description. Precisely what definite description should we use to fix the reference of "knowledge," say? Is knowledge "the intentional state that is of fundamental interest to epistemologists?" Hardly. There are lots of intentional states that are of fundamental interest to epistemologists. Is knowledge the belief state that precludes the possibility of error? Maybe. Is knowledge the true belief that results from a properly functioning belief-producing process that is operating in precisely the environment for which it was designed? Perhaps. Is knowledge the intentional state directed at truth and supported by evidence that contains no essential falsehoods? Could be. But these suggestions are precisely the controversial suggestions that we need to evaluate in our search for a plausible account of knowledge. They are not uncontroversial *preliminary* moves enabling us to begin an empirical investigation into the nature of knowledge!

The same problem affects an attempt to secure our samples through demonstrative reference. You might pick out as one of the samples of knowledge to investigate your belief that you will be alive tomorrow. Another philosopher might well claim that such a belief is hardly an instance of knowledge. We beg all sorts of interesting questions at the outset by our selection of the samples to be investigated. But it is not just the controversial character of knowledge claims that creates problems for our search for the right samples. Whatever we choose as our paradigm of knowledge might not be knowledge *even by our own lights*. Suppose, once again, you were unlucky enough to choose your samples of "knowledge" while trapped in a Gettier situation. You make such declarations as that if anything is a case of knowledge, *this* is. We start investigating the situation and discover that we relied on inferences from justifiably believed falsehoods in reaching our "paradigm" of knowledge. Letting the chips fall where they may, should we then conclude as a result of our empirical investigation that knowledge apparently involves inference from justifiably believed falsehoods? Of course not. Our empirical discovery will lead us to abandon our original declaration! And the reason is again simple. We never were willing to let an empirical investigation decide what would or wouldn't count as a case of knowledge—at least not an empirical investigation of the sort envisioned by the philosopher modeling an analysis of knowledge on the scientific "analysis" of water or lightning. When we try to understand our concept of knowledge, we think of all sorts of hypothetical situations and try to figure out the common denominators to all of the belief states that we would call knowledge in an effort to learn something about our epistemic concepts.

Sources of Knowledge and the Problem of the Criterion:

As I indicated in my introductory comments, the more familiar problem of the criterion, the problem most often associated with the work of Chisholm, concerns not the analysis of epistemic *concepts*, but the search for "sources" of knowledge and justified

belief. So as Chisholm describes the opposing views, the particularist, relying on so-called commonsense, decides that we know and are justified in believing all sorts of commonplace propositions and uses this as a *datum* to argue for the necessity of recognizing either certain sorts of justification as noninferential or for recognizing the legitimacy of various sorts of inference. The “generalist” or the “methodist,” by contrast, decides what kinds of beliefs are noninferentially justified and what kinds of inferences are legitimate, and is willing to let the chips fall where they may when it comes to implications concerning what we do and don’t know or are justified in believing. Once again, the methodology favored by the particularist seems simply to rule out the possibility that skepticism might be true. The particularist decides at the very outset that the correct epistemology will ensure that we know and are justified in believing most of what we think we know and are justified in believing. We’ll “give ourselves” whatever principles of justification we need in order to arrive at the anti-skeptical destination we are determined to reach. Most of the classical skeptics, by contrast, were methodists. Hume, for example, advanced a view about how one can legitimately move from premises to conclusions and argued from that view for skepticism.⁷

Following Chisholm, let’s illustrate the particularist’s strategy for dealing with skepticism by looking at a couple of skeptical arguments. Consider skepticism with respect to the external world. Hume famously argued that the only way to justifiably believe that there are objective, enduring physical objects, is to *infer* their existence from what we know directly about the character of fleeting subjective appearance. He further argued that the only way to get legitimately from knowledge of appearance⁸ to knowledge of objective reality is to inductively establish the occurrence of the former as evidence for the existence of the latter. But, Hume argued, we cannot correlate appearance with something other than appearance—we have no direct access to anything but appearances. We have, therefore, no justification for believing that there exist physical objects.

Or consider the problem of memory. I seem to remember having had a headache earlier today and reach the conclusion that I did have the headache. No-one my age needs to be convinced that that seeming to remember doing *X* doesn’t entail having done *X*. But once again, the prospect of inductively establishing memory as a reliable indicator of the past seems hopeless. Any inductive argument would rely on *past* correlations of memory “experiences” with the events we took them to be signs of, and without relying on the legitimacy of the very memory under skeptical attack, we would have no access to such past correlations. The radical empiricist seeking to identify foundational justification with internally accessible infallible justification, and committed to the view that the only sort of nondeductive inference that is legitimate is inductive inference, will inevitably be led to a radical skepticism about the past.

⁷ It’s also possible to argue for a kind of skepticism from the premise that the problem of the criterion is insoluble. The idea would be that there simply is no way of reaching conclusions about particular cases of what we know without knowledge of methods, and there is no way of arriving at knowledge of methods without the guidance provided by particular claims about what we know. So we can never be in a position to know what we know (even if we have first-level knowledge).

⁸ Chisholm would prefer “self-presenting states

The particularist with whom Chisholm sympathizes responds that if Hume is right and we cannot use either deduction or enumerative induction to legitimately infer the existence of physical objects or to reach conclusions about the past, then there must be some other source of justified belief. It is simply absurd, the particularist argues, to deny that we have the knowledge or justified belief under skeptical attack. If we aren't going to use *modus tollens* here to reject a philosophical position, we'll never use it. Chisholm, for example, agrees with Hume that we aren't getting where we want to go relying exclusively on deduction and induction from infallibly believed foundations, and is willing to infer from this that we had better recognize additional epistemic principles that will sanction the beliefs under skeptical attack. So in the 3rd edition of *Theory of Knowledge*, for example, Chisholm advances the following principles, principles tailor-made to counter the skeptic:

For any subject *S*, if *S* believes, without grounds for doubt, that he is perceiving something to be *F*, then it is beyond reasonable doubt for *S* that he perceives something to be *F*. (p. 76)

For any subject *S*, if *S* believes without grounds for doubt, that he remembers perceiving something to be *F* (where *F* is a sensible characteristic) then it is beyond reasonable doubt for *S* that he does remember perceiving something to be *F*. (p. 81)

Particularism, Methodism, or Reflective Equilibrium:

Epistemological Internalism and the Problem of the Criterion:

So who has the correct approach to the philosophical investigation of sources of knowledge and justification? Is it the particularist or the methodist? Or should we reject both as extreme positions and sign on to the method of reflective equilibrium, an approach that recommends that we try to maximize coherence between our "intuitions" about the legitimacy of methods and our "intuitions" about particular cases of knowledge and justified belief?

It is important, I think, that however one answers these questions one does so in a principled way, a way consistent with one's other epistemological views. Chisholm, for example, seems to me to be in a precarious position. According to Chisholm, coherence among one's beliefs can increase the level of epistemic justification those beliefs enjoy. But he is, in the final analysis, a confirmed foundationalist. There are beliefs that enjoy foundational justification, and all other beliefs owe their justification in part to these foundationally justified beliefs. Furthermore, like most classical foundationalists, Chisholm includes among the candidates for foundational justification, necessary truths knowable *a priori*. And that brings us back to epistemic principles. For Chisholm, epistemic principles describe conditions under which a belief enjoys a certain positive epistemic status (the terms of epistemic appraisal are all defined in terms of the primitive "being more reasonable to believe than"). So what is the modal status of an epistemic principle and what is the nature of our justification for believing an epistemic principle?

It is clear from the corpus of Chisholm's work that he doesn't take an epistemic principle to describe *contingent* circumstances under which beliefs enjoy an epistemic

status. He takes epistemic principles to state necessary truths, and presumably allows that one can know such truths *a priori*. If there is any doubt about this, one need only reflect on Chisholm's commitment to access internalism. Chisholm holds one always has the ability to directly access the epistemic status of one's beliefs (1977, 116).⁹ And that's because he thinks that the epistemic status of beliefs depends on subjective states that are self-presenting, and the truth of epistemic principles that one can know *a priori*.

So what business does Chisholm have embracing the particularist's position? He should be the paradigm of a philosopher who endorses the claim that one can acquire unproblematic knowledge of epistemic principles and employ such knowledge in reaching conclusions about what we know and are justified in believing. He really should let the chips fall where they may. Of course, you might suppose that the question of whether Chisholm can consistently endorse particularism is at best a matter of historical interest. But my observations about what Chisholm is committed to extend to any philosopher who embraces a similar sort of access internalism. I have argued elsewhere (1996, 2004) that one cannot consistently embrace plausible internalist access requirements unless one takes the principles that sanction inference to be synthetic necessary truths knowable *a priori*.¹⁰ If one accepts, for example, the view that I call *inferential internalism* and argues that to justifiably infer *P* from *E* one needs justification for believing that *E* makes probable *P*, the only way to avoid a vicious regress is to acknowledge that there are some propositions of the form *E* makes probable *P* that are knowable *a priori*. And I don't see how one can take that position without rejecting the view that one should hold one's conclusions about epistemic principles hostage to conclusions about particular claims to knowledge and justified belief.

In arguing for this position I want to emphasize again a point made earlier in connection with the particularist/generalist controversy over the nature of our access to the truth of philosophical analyses. It is doubtless true that in trying to discover plausible epistemic principles one might well consider particular *inferences* (in contrast to particular claims about what we know or are justified in believing). Think about how you try to teach undergraduates the rules of deductive inference. To be sure one could just trot out general principles like *modus ponens* and *modus tollens* and tell the students to think carefully about the principles until they start seeming pretty obvious. And this might work. But it often helps to ask students to think about particular applications of *modus ponens* and *modus tollens* as a heuristic device. The fallacy of affirming the consequent might not seem like a fallacy until one thinks about inferring that Chisholm was a dolphin from the fact that if he were a dolphin he would have lungs, and he does have lungs. But this isn't an argument for Chisholmian particularism. We are not trying to reach conclusions about principles of reasoning from particular claims about what we know or are justified in believing. We are trying to get the students to discover *a priori* the legitimacy of a particular inference in an effort to get them to see *a priori* the more general principle of reasoning instantiated in this particular inference. Our discovery of the general principle is arrived at through consideration of the particular inference only in the sense in which our original "seeing" that two plus two equals four was aided by our

⁹ The principle here is advanced concerning knowledge, but it is clear that Chisholm would apply it to all other epistemic properties of one's beliefs.

¹⁰ I think that this is one of the main thrusts of Bonjour's (1998) defense of the *a priori*.

kindergarten teacher's illustrating that *a priori* truth with the truth about the four apples arranged on the desk.

Epistemological Externalism and the Problem of the Criterion:

So I've argued that most classical internalists, ironically the very kind of internalists of whom Chisholm was a paradigm, should place themselves firmly in the generalist's camp. But what of the increasingly popular externalisms embraced by many contemporary epistemologists? The vast majority of externalists reject the claim that all epistemic principles are necessary truths.¹¹ The reliabilist, for example, is clearly committed to the view that once we leave deduction the legitimacy of an inference is a matter of contingent fact. Should such a philosopher embrace particularism or methodism?

For me, the question is difficult to answer because I have never been sure how externalists do decide which beliefs are or are not justified, or why they place such confidence in their "intuitions" about which particular beliefs are justified. Again, let's take as our paradigmatic externalist, the reliabilist. We certainly won't get very far trying to discover where their sympathies lie by looking at their practice. As far as I can tell most reliabilists do seem to take for granted that most of what they take to be justified beliefs are in fact justified. And they do sometimes seem to be guided by these "intuitions" when it comes to addressing problems raised both for their analysis of justification and for their views about what should count as a reliable belief-producing process. In this respect they look a bit like the particularist Chisholm describes. On the other hand, reliabilists seem to me to have equally confident, unswerving faith in the reliability of commonly used process types like perception, memory, induction, and the like. Again, it's just not clear to me what they think their epistemic license is to place such faith in the reliability of the various belief-producing mechanisms they employ.

Rather than look to the externalist's practice for insight, let's instead try to figure out how a consistent externalist, a reliabilist, for example, should go about reaching conclusions about the epistemic status of particular beliefs, and the legitimacy of various sources of justification.¹² Consider memory. As I said earlier, most reliabilists I have read are quite certain that most of what they seem to remember happened, that most of their beliefs based on an apparent memory are justified, and that memory is a more or less reliable source of information about the past. To be sure, any careful discussion of the issues would require the reliabilist to address head on now familiar problems facing their account of justification. In particular, before one assesses either the justificatory status of a memory-produced belief, or the general reliability of memory, it would be critical to address the well-known generality problem. Empirical research indicates that there are radically different kinds of memory, and one sort may be much more reliable than another. Furthermore, commonsense itself suggests that we may want to treat as a different process kind memory-induced belief in old people and memory induced beliefs in younger people. For all I know the reliabilist will end up with memory after an all

¹¹ Even externalists will allow that some belief-producing processes are necessarily truth-conducive—processes that realize the rules of deductive inference, for example.

¹² It what follows I'm going to use the reliabilist as the paradigm of an externalist in epistemology. I believe that all of what I say will apply *mutatis mutandis* to other versions of externalism.

night bender being a different belief-producing process from memory after taking a course on how to improve your memory. Thankfully, I'm not a reliabilist so I don't have to worry about how to divvy up possible belief-producing processes about the past. There is the further complication that verbs like "remember" probably are usually used to report factive states.¹³ If my belief that I put the keys on the table results from remembering that I did in this factive sense, then, trivially, the belief-producing process is 100% reliable. For any interesting epistemological discussion of memory, we will need to identify memory-induced beliefs as those whose input is something more like seeming to remember, where seeming to remember is perfectly compatible with the apparent memory being nonveridical.

With all that as preliminary, let us suppose that we've identified a belief-producing process that takes as its input something like seeming to remember having done *X* and has its output the belief that *X* occurred. We ask the reliabilist to justify both his belief that the belief that *X* occurred is justified and his belief that the input/output mechanism that results in that belief is reliable—is a legitimate source of justification. How should the reliabilist proceed? Well, I suppose he might try an inductive argument. He might claim to remember seeming to remember a great many things where most of the time the apparent memory turned out to be accurate. I seemed to remember seeming to remember having put the keys on the table, and the keys were there. I seemed to remember having had cornflakes for breakfast and I did. Get enough of these correlations and we're in business generalizing to the reliability of apparent memories.

One doesn't have to be in the clutches of skepticism to wonder whether the above argument isn't suspect. To inductively infer that the belief-producing process in question is reliable, we need justified beliefs in the premises of our argument. According to reliabilism, we may be justified in believing that the keys were on the table, and we may be justified in believing that we seemed to remember having put the keys on the table, but only if the relevant sort of memory upon which we rely is reliable. But that is precisely what we are trying to establish. Our belief in the conclusion of our argument based on our acceptance of the premises is only as justified as our belief in the premises, and that justification presupposes that the conclusion is true. I have argued elsewhere (1996) that none of this should bother the consistent reliabilist. The facts about inputs and reliabilities that determine the epistemic status of our beliefs, according to the reliabilist, are independent of our access to those facts. Given the view, the justification of first-level beliefs does not require any beliefs about the reliability of methods by which the beliefs were acquired. It's enough that the world co-operate so as to supply the relevant correlations. And, of course, if the metaepistemological view were correct, the matter should be no different when it comes to the epistemic status of second-level beliefs about the epistemic status of first-level beliefs. It's simply not necessary that we have any sort of epistemic access to the legitimacy of whatever methods we employ in reaching our conclusions about the epistemic status of first-level beliefs.

Again, I've argued in a number of places that the sort of epistemic justification in which philosophers are interested is tied to getting *assurance* of truth, and I don't see how satisfying the reliabilist's conditions for epistemic justification gets oneself

¹³ A psychological state directed at an object is factive if the truth of the statement describing the state entails that the object exists. So my seeing a table is factive in that my doing so (arguably) requires the existence of the table. My fearing ghosts isn't, because the fear can exist without the ghosts.

assurance of anything. If I start worrying that my memory is getting very bad, I don't think I can assure myself that it's really not so bad by apparently remembering all sorts of recent successful employment of memory (though I can, perhaps, confirm my fears by seeming to remember all sorts of unsuccessful reliance on memory). But here, I'm interested primarily in the issue of where the externalist, in this case represented by the reliabilist, should stand on the particularist/generalist controversy concerning the problem of the criterion

The upshot of the above discussion suggests initially that the reliabilist is *some* sort of generalist/methodist. Again, given the reliabilist's understanding of epistemic justification, one must discover something about the source of a belief to reach a conclusion about its epistemic status. To be sure, the track argument discussed above employs premises that describe particular instances of memory and particular conclusions about the past. But it does *not* involve particular claims about the epistemic status of our beliefs in these propositions. To be sure, without justification for believing the premises of the argument, we have no justification for believing the conclusion of the argument, but that doesn't alter the fact that the premises make no claim about the epistemic status of particular beliefs. Chisholm's particularist employs premises describing the epistemic status of various beliefs in order to reach conclusions about the legitimacy of epistemic principles.

But if the reliabilist is a kind of generalist, it is important to recognize the difference between this sort of generalist and the sort that Chisholm had in mind. Given reliabilism, it is *trivially* true that every conclusion about the epistemic status of a belief involves a conclusion about the "legitimacy" of its source. But only because a conclusion about the epistemic status of a belief just *is* a conclusion about its having the right kind of source.¹⁴ That position does not imply that one reaches conclusions about the epistemic status of beliefs through prior or independent knowledge of the legitimacy of epistemic principles. The consistent reliabilist will insist that one needn't know much of anything about the *specific* nature of the belief-producing mechanism that leads to a belief in order to reach a justified conclusion about that belief's epistemic status. As I have pointed out many times, the reliabilist places no *a priori* restrictions on what beliefs can be justified, even justified noninferentially. I might be so constituted that when I entertain the proposition that a belief is justified I instinctively respond with either an endorsement or a rejection of its having positive epistemic status. And I might be almost always right. I might be the epistemic counterpart of the infamous "chicken-sexer" much discussed in the literature.¹⁵ If I am, and if the reliabilist is correct, then I can come to know that various beliefs are justified and I can do so without even understanding *how* I am reaching these conclusions. If the reliabilist is right a conclusion about the epistemic status of a belief *is* equivalent in meaning to a conclusion about the belief's having the right sort of cause, but reaching the conclusion doesn't involve an examination or investigation of that cause. If the *robust* generalist/methodist thinks that one can only reach legitimate conclusions about the epistemic status of beliefs by *first* finding out

¹⁴ In this sense, the classic internalist might also be trivially committed to a form of generalism when it comes to reaching conclusions about inferential justification. The classic recursive analysis of justification might define inferential justification in terms of legitimate inference from noninferentially justified belief.

¹⁵ The chicken-sexer is supposed to have the uncanny ability to look at a chick and predict its sex even though (according to the story) he can't tell how he arrives at the conclusion.

which specific epistemic principles are legitimate, then the reliabilist (and more generally the externalist) is neither a robust generalist nor a particularist. As far as I can see, on this controversy, the externalist doesn't have a dog in the hunt.

Reflective Equilibrium and the Problem of the Criterion

In the preceding discussion, the method of reflective equilibrium has been conspicuous by its absence. And given the widespread use of the method, this might seem a bit odd. In the final analysis, however, it seems to me that one can deal with the method of reflective equilibrium in short order. Whether one is searching for a correct analysis, or trying to discover the sources of knowledge and justified belief, one can legitimately employ the method of reflective equilibrium only if the coherence theory of justification is true. And it is not. Let me explain.

According to proponents of bringing beliefs into reflective equilibrium, one takes beliefs one has about particular beliefs being justified, one takes one's beliefs about various epistemic principles being justified or unjustified, and one makes whatever adjustments will allow one to maximize retention of beliefs. When the smoke clears and the dust settles, however, what could possibly justify this approach to arriving at the truth of the matter but the idea that justification is a function of coherence? The beliefs that survive the winnowing process are justified, presumably, in virtue of occupying a place in a system of beliefs that coheres. We haven't the space here to engage in a full-scale evaluation of the coherence theory of justification. I have argued elsewhere, however, that a coherence theory comes in either an externalist or an internalist version. On the externalist version of the view, my beliefs are justified provided that they do in fact cohere (in whatever sense of coherence the coherence theorist develops) and I don't need any access to the fact that the beliefs do cohere. This version of the coherence theory is vulnerable to devastating counterexamples. One need only imagine a person who whimsically chooses to believe a great many propositions based on bizarre hunches and who by sheer happenstance ends up believing propositions that mutually support each other in ways that only a master logician could discover. Such a person has coherent beliefs but also has wildly irrational beliefs.

If in the face of this objection one turns to an internalist coherence theory, and allows that one must have access to coherence among one's beliefs in order for those beliefs to be justified, one faces the regress argument that Bonjour (1985) made famous. One's access to coherence will involve access to one's beliefs states, an access that is either foundational (and thus inconsistent with the coherentist's position) or is dependent on relations of coherence to which one needs access—the beginning of the regress.

In short I'm not interested in whether someone rests comfortably with a belief system in reflective equilibrium, regardless of the subject matter of those beliefs. I've known many a philosopher, and the odd paranoid schizophrenic, with wonderfully coherent belief systems where I'm quite convinced that the beliefs are mostly false and mostly irrational.

Richard Fumerton
University of Iowa

References

- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- _____. 1998. *In Defense of Pure Reason* (London: Cambridge University Press).
- Chisholm, R. M. 1977. *Theory of Knowledge, 3rd Edition*. Prentice-Hall.
- Fumerton, Richard. 1983. "The Paradox of Analysis," *Philosophy and Phenomenological Research*, 43, 477-497.
- _____. 1989. "Russelling Causal Theories of Reference." In *Rereading Russell*, eds. Savage and Anderson. Minneapolis: University of Minnesota Press, 108-18.
- _____. 1996. *Metaepistemology and Skepticism*. Boston: Rowman and Littlefield.
- _____. 2004. "Epistemic Probability." *Philosophical Issues*. Vol. 14, 149-64.
- Gettier, Edmund. 1963. "Is Knowledge Justified True Belief?" *Analysis*, 23, 121-3.
- Kornblith, Hilarly. 2002. *Knowledge and Its Place in Nature*. Oxford: Oxford University Press.
- Russell, Bertrand, 1948. *Human Knowledge: Its Scope and Limits*. New York: Simon and Schuster.