

Trustworthiness

Catherine Z. Elgin

Abstract: I argue that trustworthiness is an epistemic desideratum. It does not reduce to justified or reliable true belief, but figures in the reason why justified or reliable true beliefs are often valuable. Such can be precarious. If a belief's being justified requires that the evidence be just as we take it to be, then if we are off even by a little, the belief is unwarranted. Similarly for reliability. Although it satisfies the definition of knowledge, such a belief is not trustworthy. We ought not use it as a basis for inference or action and ought not give others to believe it. The trustworthiness of a belief, I urge, depends on its being backed by reasons – considerations that other members of the appropriate epistemic community cannot reasonably reject. Trustworthiness is intersubjective. It both depends on and contributes to the evolving cognitive values of an epistemic community.

During their formative years, many epistemologists seem to have obsessed over grades. They maintain that our overriding epistemic objective is to believe as many truths as possible and to disbelieve as many falsehoods as possible, effectively construing cognitive life as a giant true/false test. On this construal, the rationale for adding justification or reliability to the true belief requirement on knowledge is prudential. We improve our odds of believing truths and disbelieving falsehoods by restricting our beliefs to those that are justified or reliable. But life is not a true/false test. Truth value does not determine cognitive value. So to restrict the focus of epistemology to factors that would maximize our prospects of acing the test is unwise. It substitutes a thin conception of knowledge for thicker conceptions of epistemic states that are more valuable.

One thickening agent is trustworthiness. Not all true beliefs are trustworthy. Lucky guesses are not. Neither are beliefs with shaky supports. If a true belief is trustworthy, one can safely and responsibly act on it, use it as a premise in assertoric

inferences, and give others to believe it. If it is untrustworthy, one cannot. The cognitive value of a true belief is to a considerable extent a function of its trustworthiness.

Thick and Thin

That a political ploy instantiates the predicate ‘sleazy’ conveys information both about what sort of maneuver it is and about how it is to be evaluated. ‘Sleazy’ is a thick term, being simultaneously descriptive and evaluative. ‘Good’ and ‘red’ are thin. That an item instantiates the predicate ‘good’ reveals nothing about what sort of thing it is; that it instantiates ‘red’ reveals nothing about whether it is any good.

The ‘thick or thin’ distinction between concepts or terms is typically framed in a way that suggests that determining a concept’s girth is a job for conceptual analysis. If so, to understand what makes ‘sleazy’ a thick term is to properly analyze the word ‘sleazy’. My project is not conceptual analysis. I do not care how norm-laden terms are defined, or whether they admit of definition. I am interested in the constitution and configuration of the epistemic realm. I want to know what elements comprise it, how they relate to each other, and where and how the boundaries between them are drawn. Epistemology, as I understand it, should produce a theory, not a dictionary.

Nevertheless, we need to begin with a semantic point. Thick terms like ‘sleazy’ are not mere compounds that can be factored into descriptive and evaluative components. By way of contrast, consider the newly coined predicate ‘gred’, which applies to all and only things that are both good and red. Such a predicate is both descriptive and evaluative. A mitten satisfies the descriptive requirement (by being red) and satisfies the evaluative requirement (by being good). But a gred item’s being red has no bearing on its being good; nor has its being good any bearing on its being red. The mitten qualifies as

red because of its color; it qualifies as good because of its capacity to warm a hand. The mitten just happens to instantiate the two predicates; it belongs to their intersection. Because we can prize apart the descriptive and the evaluative components, ‘gred’ is not a genuinely thick term. It is simply a contraction.

The relation between the descriptive and the evaluative elements in thick terms like ‘sleazy’, ‘truthful’ and ‘trustworthy’ is tighter. Thick terms involve a fusion of descriptive and normative elements. They do not admit of the sort of factor analysis that ‘gred’ does. Although ‘sleazy’ is a descriptive predicate whose instances we readily recognize, there is no evaluatively neutral way to delineate its extension. We have no way to identify the class of sleazy politicians without intimating that they merit disapproval. If epistemic terms are similar, then we should have no way to identify their extensions without intimating that they are worthy of approval (or disapproval) on cognitive grounds.

In *Truth and Truthfulness*, Williams (2002) shows how the thick concept of truthfulness depends on and diverges from its thin descriptive precursor – the concept of uttering truths. I briefly review his discussion because it illustrates how thick terms relate to their less corpulent ancestors. As a first approximation, a truthful person is someone who utters or is disposed to utter truths. We can identify such people without taking a stand on the value of uttering truths, just as we can identify terse people without taking a stand on the value of brevity. The evaluative element enters the picture later, when we deem uttering or being disposed to utter truths a good thing. At the first step, there is no fusion. The descriptive and evaluative elements can easily be separated. But, Williams notes, the virtue of truthfulness does not consist exactly in uttering or being disposed to

utter truths. It involves both more and less than that.

Obviously, we do not expect the truthful person to utter or be disposed to utter only truths. Everyone makes mistakes. So if anyone is to count as truthful, truthfulness cannot require uttering only truths. It might, however, require saying only what one thinks is true. That would leave room for mistakes but exclude lies. This is a step in the right direction, but it demands too little for at least two reasons.

First, misleading is as much a violation of truthfulness as lying is. It is possible to utter only truths and, by exploiting implicatures, induce one's audience to form a false belief. If I tell my editor, 'I can't get the paper in before next Tuesday,' I give her to believe that I can get it in on Tuesday. If I know there is no chance it will be done until Friday, I have not been truthful, even though the sentence I uttered was true.

Second, being a virtue, truthfulness involves responsibility. We would not consider a person truthful if she were unduly careless or gullible in her belief formation. Suppose, despite her excellent education, Joan believes everything she reads in the tabloids and indiscriminately reports her beliefs. She believes with equal conviction that space aliens landed in Detroit and that Barry Bonds took steroids. If she uttered the true sentence 'Barry Bonds took steroids,' we would not consider her truthful. For she was equally prepared to utter the falsehood, 'Space aliens landed in Detroit.' Her utterances are not credible, since she does not adequately filter out falsehoods. Although she says what she believes, she lacks the virtue of truthfulness.

Nor is it enough to impose a responsible belief requirement. Truthfulness does not require uttering every truth one responsibly believes. Some truths are private. When asked about them, Williams believes, evasions, obfuscations, even lies are permissible.

Your questioner may have no right to know whom you plan to vote for or how much you weigh. If refusing to answer unduly intrusive questions is not feasible, and ‘It’s none of your business!’ is too blunt, Williams maintains, a truthful person need not tell the truth. Her interlocutor should recognize that in asking such questions, she oversteps the bounds within which truthfulness is required.

Truthfulness, as Williams characterizes it, is intimately related to, but does not supervene on, uttering truths. For the requirements on speaking truthfully diverge from those on uttering truths. The reason for the divergence is pragmatic. Because a more nuanced attitude towards truths is more valuable than uttering all and only things one believes to be true, that more nuanced attitude is the one we should admire and cultivate. ‘Truthfulness’ does not admit of factor analysis because, without appealing to evaluative considerations, we have no way to delineate the extension of the class of topics about which one should speak the truth.

If Williams is right about the relation of a thick concept to its thin descriptive precursor, then thick concepts are not a mere convenience. Without such normatively loaded concepts, we could not partition the world as we do. So the idea that thick concepts are simply contractions, that we could make do with morally thin concepts – good, bad, right, and wrong – plus our stock of purely descriptive concepts, is incorrect. Without the thick concepts, we would have no way to mark out the extensions in question. If epistemic concepts are thick, we should expect them to display the same pattern. They should be anchored in, but not supervene on, a descriptive core. They should mark out extensions that we have no evaluatively neutral way to demarcate.

Trustworthiness

A critical factor in fixing the contours of ‘truthfulness’ is the idea that truthful people are trustworthy. Within the realm where truthfulness is required, when a truthful speaker says that p , we can take his word for it. So the epistemological question is whose word should we take? Which informants are trustworthy?

An anorexic answer is that we should trust anyone who speaks the truth. Whatever his background, character, or track record, if we want to know whether p , and he speaks truly when he says that p , we get what we want when we take his word. This is far too thin to be remotely plausible. It would advocate taking the blind man’s word for the route to Larisa. It would even advocate taking the word of a blind liar who has never heard of Larisa, if he happened to be right. But if we thought our informant was incompetent or insincere with respect to p , we would be epistemically irresponsible to take his word. Similarly for other sources. If we have no reason to trust our evidence or the reliability of a perceptual deliverance, we would be idiots to trust it. Evidently, something is trustworthy only if it would be epistemically responsible for us to count on it.

This is where epistemologists introduce the requirement that for a belief or utterance to be epistemically estimable its truth must be, in an appropriate sense, nonaccidental. Such a belief must be suitably related to something (perhaps a method, perhaps evidence) that is truth conducive. This captures the idea that lucky guesses do not count as knowledge. Evidently the descriptive core of the epistemic concept of trustworthiness is truth conduciveness. As a first approximation then, for something to be trustworthy its obtaining must be truth conducive.

Reliabilists maintain that a belief is relevantly nonaccidental if it is caused or

sustained by a reliable mechanism. We have reason to trust the sighted guide to Larisa but not the blind guide because vision is a reliable mechanism for providing information about routes. Reliable mechanisms, they maintain, are truth conducive.

There are at least two difficulties with this position. The first is brought out by a variant on Bonjour's example (1985, 38-40). Suppose Marie has ESP. She is subject to reliable extra sensory deliverances. In 93.7% of the cases where she has such a deliverance, it is true. She is also subject to the normal sorts of hunches and intimations, and her track record with these is no better than anyone else's. Moreover, the deliverances of ESP are phenomenologically no different from hunches. Nothing in her experience enables her to tell whether a premonition that pops into her mind is a deliverance of ESP or a mere hunch. Nevertheless, the etiology is different. The deliverances of ESP result from a reliable perceptual mechanism; the hunches result from unreliable mechanisms. (Possibly an fMRI would reveal the difference, if only we knew where to look.) Moreover, Marie knows that scientists believe that ESP does not exist. So although her genuine extra sensory perceptions are reliable, she has no reason to trust them. She has no way to distinguish them from her unreliable hunches, and has good reason to suspect that there is no such thing as ESP. Evidently, mere reliability is not enough. Minimally, the subject should have grounds for believing that a process is reliable.

Even that seems inadequate. The second worry brings this out. Some therapeutic regimens work for no known reason. An effective treatment for bladder cancer involves flooding tumors with live bacteria. Somehow that kills the cancer cells. The statistical evidence is strong; the therapy is effective. But patients feel epistemically insecure about

trusting their lives to a course of treatment whose effectiveness cannot currently be explained. This suggests that even knowing *that* a process is reliable is not enough. We want to understand (at least roughly) why and how it works.

This might seem to push in the direction of evidentialism. But the same problem arises for the evidentialist. It is not enough to have sufficient evidence, we need adequate grounds for believing that we have sufficient evidence. Perhaps in the vast and motley collection of data, there is sufficient evidence to explain why the cancer treatment works. But scientists have not yet found a way to identify and extract the relevant evidence. So the success of the therapy remains a mystery. As things stand, the relation between the truth and the evidence is epistemically inert. What we need is not just claims we have sufficient evidence for, but claims we have good grounds for believing that we have sufficient evidence for, where grounds are factors that enhance objective probability.

Unfortunately, it seems not to be enough that our grounds be good; we want some assurance that they are good. A regress looms. If we need sufficient evidence that we have sufficient evidence, and sufficient evidence that we have sufficient evidence that we have sufficient evidence, . . . , our situation looks bleak. Similarly if we have to demonstrate the reliability of our methods for assessing reliability, and demonstrate the reliability of those methods, and so on without end. To block such a regress, the requisite assurance cannot consist entirely of grounds for our grounds.

Before considering how anything other than objective probability enhancers could help here, let us consider what besides truth conduciveness might be epistemically desirable. By identifying and justifying other epistemic desiderata, we put ourselves in a better position to resolve our dilemma.

Solidity

Let us call the considerations that are taken to constitute its grounds the *supports* for a belief. Although a belief's grounds are the factors that affect its objective probability, assessment of grounds is a normative matter. It depends on our epistemic standards, cognitive methods, and background beliefs. A true belief is justified if it is in fact adequately backed by evidence; it is reliable if it is secured by a method that is in fact reliable. Such a belief might nevertheless be vulnerable. If its truth conduces and its relation to them must be almost precisely as they are for the belief to be justified, then any small perturbation undermines its epistemic standing. Such a belief is shaky. *Shakiness* and its antithesis, *solidity*, are multifaceted. Supports are solid if they are secure, stable, and robust. Otherwise they are shaky. Each of the facets is a matter of degree. The question is not whether a belief can withstand losses or shifts in support, but how resistant to such changes and how resilient in the face of such changes it is. That is, how much perturbation it can withstand.

A belief is *secure* to the extent that its supports are unlikely to shift. In that case, the considerations currently regarded as telling in its favor are apt to continue to be so regarded. A belief is *precarious* to the extent that its the supports are likely to shift. Supports shift when what has been considered evidence loses its status as evidence and/or when what has been considered not to be evidence acquires the status of evidence. They also shift when mechanisms that had been considered reliable come to be considered unreliable and/or when mechanisms that had been considered unreliable come to be considered reliable. A recently developed lie detector test correlates particular patterns of neural activation with lying. The correlation may in fact hold, in which case the test is

reliable. If so, it provides good evidence that a subject is (or that he is not) lying. But in view of how spotty and provisional our knowledge of the relevant neurology is, the test and its results are insecure. It would not be surprising if the test is unreliable. Perhaps other behaviors activate those neural circuits, or in skilled liars they are not activated. Here insecurity does not stem from unreliability. For all we now know, the test is reliable. But in our current epistemic circumstances, we should not be confident that it is reliable. As things stand, we would be epistemically irresponsible to place much trust in it. A diagnosis of a bacterial infection is far more secure. Given the depth and breadth of our understanding of bacteria, antibodies, and blood tests, the evidence of antibodies that a blood test reveals is likely to continue to be considered a reliable indication of bacterial infection. In the case of the blood test, we are in a position to be confident that our evidence is good and our method reliable.

A belief is *robust* to the extent that its being justified can survive revisions in support. It is *fragile* to the extent that its justification or reliability is undermined by such revisions. Counting ticket stubs is a good way to determine how many people attended a concert. Suppose a careful count yields the number 2603. Fred's belief that 2603 people attended the concert is supported. But his support is fragile. It would be unsurprising if a recount yielded a slightly different number. On the basis of the original count, Flo believes that about 2600 people attended the concert. Her belief is more robust, since, given that the first count was careful, any reasonable recount can be expected to yield a number in the neighborhood of 2600. A justified belief can be fragile. If the support is adequate and turns out to need no revision, then the belief remains justified. But it seems risky to place much trust in a belief whose justification could easily be lost.

A belief is *stable* to the extent that it can withstand some loss of support. To the extent that its acceptability is threatened by such losses, it is *unstable*. Support is lost when what had once been thought to tell in favor of a belief is no longer thought to do so. If the only basis for believing that Willy robbed the bank is the eyewitness testimony of one bank teller, the belief is unstable. Should that witness turn out to be severely myopic, she is unreliable and the evidential value of her testimony is null. But if there are multiple sources of information – several independent witnesses, security cameras, and fingerprints – then the discovery that any one of them is unreliable leaves us with ample support for the belief that Willy robbed the bank. Even if its support is unstable, a belief is justified so long as the grounds are adequate. If the one eyewitness has good eyesight, saw the robber in good light, and was not flustered by the robbery, her testimony may be sufficient to justify the belief that Willy robbed the bank. But since the belief is based on a single source, it is unstable. An unstable belief is epistemically undesirable, because with respect to it, we are vulnerable.

A shaky belief is epistemically unsatisfactory even though it satisfies the standard conditions for knowledge. It is a justified or reliable true belief that is not a Gettier case. But given the vicissitudes of epistemic life, it could easily be wrong. Granted, a solid belief could also be wrong. But we would have to be wrong about a lot and/or off by a lot to be wrong about it. Where support is shaky, if we are off even by a little, the belief is unjustified. Solidity then underwrites trustworthiness. If a belief rests on solid supports, we can reasonably and responsibly act on it, can use it as a premise in reasoning, and can in good conscience convey it to others.

Perhaps solidity is too much to demand. One alternative is to insist that ‘When

you're right, you're right'. Knowledge requires adequate grounds. Whether the support is shaky makes no difference. This stance has the virtue of not setting us off on an interminable search for grounds for our grounds for our grounds. It allows us to know a little without having to know a lot. It concedes that children and non-human animals can know, even though they lack the resources to form or justify second order beliefs. But it does not do justice to our epistemic ambivalence about shaky beliefs.

A second response is to incorporate solidity into the requirements for knowledge, recognizing that more is needed for justification or reliability than we might have thought. But the considerations that bear on solidity are of a higher order than evidence or reliability. Providing reason to believe that our grounds for believing that p are solid is different from simply providing additional grounds for p . It is different from providing a greater measure of reliability than our previously accepted standard of reliability for propositions such as p . We can augment our evidence that the lake is polluted with PCPs by testing more samples from the lake. We can recalibrate our measuring device so that it yields fewer false positives, thereby increasing its reliability. Both improvements are different from validating our method for testing for PCPs. But if we insist that knowledge requires validating our methods, and validating the methods for validating our methods, and so forth without end, the epistemic enterprise looks hopeless.

Despite the dismal appearances, we need not choose between having nothing against shaky grounds and embarking on an endless quest for a rock solid basis. It helps to ask why we want solid support. What is wrong with a shaky belief, assuming it is grounded? If our support is shaky, if we could all too easily be wrong, we ought not be confident in our belief. It is unwise to use it as a basis for further reasoning or for action,

and it is irresponsible to give others to believe that they can safely do so. If I take it that I cannot safely rely on p , I cannot in good conscience give you to believe that p . So perhaps we can make some headway, if we reverse the question and ask when it is *not* irresponsible to convey my belief to someone else without qualification. The short answer seems to be that I can responsibly convey my belief to you when I have good enough reasons for my belief. If I am in a position to give you to believe that p , I should be in a position to give you reasons why you should believe that p , and why you can safely and confidently act on p . This requires more than assuring you that my evidence is sound, or that my methods are reliable, or even that my beliefs are also solidly grounded. I should be in a position to assure you that if you had the requisite background, intelligence, and motivation, you would find my evidence sound, my methods reliable, and my relevant beliefs solidly grounded. I should not only reflectively endorse my belief, but be in a position to assure you that to the extent that your interests are cognitive, you too should reflectively endorse it. Indeed, I should not reflectively endorse my own belief unless I believe that, if they had such grounds, other similarly situated epistemic agents would be in a position to reflectively endorse it as well.

Reasons

Grounds are objective probability enhancers. Reasons are considerations that *can be adduced to support a conclusion*. Epistemically inaccessible factors then are not reasons. Nor are all epistemically accessible factors. The question is what epistemically accessible considerations can properly be adduced to support a conclusion. Drawing on Scanlon (1998), I suggest that a consideration c can be adduced to support p only if others who are competent to assess the relation between p and c and motivated to know

whether p could not responsibly reject the claim that c supports p .

Reasons are public. No purely subjective consideration counts as a reason; for others who are similarly motivated to know whether p could responsibly reject private intimations. However heartfelt, Pat's feeling that the Red Sox will win is not a reason to think that they will win. It is merely a hunch. Moreover, since nothing epistemically inaccessible qualifies as a reason, a reliable method for establishing that p is not a reason to accept or believe that p until we establish that the method is, or is likely to be, reliable. Currently, then, there is no reason to accept the deliverances of Marie's ESP. And a motley collection of evidence does not supply a reason to believe that p , unless we see how it or its members bear on the truth of p . Even if the census data contains evidence that people born on a Tuesday are less likely than other Americans to live on streets named for trees, merely having the census data affords no reason to believe it. To constitute a reason, the relevant evidence would have to be extracted.

Nor does any truth that fails to satisfy our standards of relevance count as a reason. Considerations that enhance the objective probability that p are not reasons if according to current standards of relevance, they have no bearing on whether p . Perhaps, left-handed people have an increased probability of developing arthritis. Still, so long as we are unaware of the correlation, Ben's being left-handed affords no reason to think that he is especially at risk for arthritis. The correlation is epistemically inert. Moreover, supports that are excessively shaky are apt not to count as reasons. If arriving at the conclusion that p requires using exquisitely calibrated instruments, making extraordinarily exacting measurements, and performing enormously complicated calculations, others who are competent with respect to the subject and motivated to know

whether p could responsibly consider the support inadequate. Something in the sequence of considerations adduced to support the conclusion could easily be slightly off; and if anything is even slightly off, the conclusion is unwarranted.

Reasons are keyed to epistemic circumstances. Public standards of evidence and relevance rest on background assumptions about the topic under investigation and what is known or reasonably believed about it. Considerations that at one point in history could not reasonably be rejected may be readily rejected later, when matters are better understood. In 1975, ‘stress causes ulcers’ was widely accepted. At that time, the contention that the high incidence of ulcers among Wall Street traders was due to their stressful jobs could not reasonably have been rejected. Once the bacterial basis for ulcers was discovered, stress can no longer be adduced as a reason why so many traders have ulcers. A factor that functions as a reason in one cognitive environment may fail to so function in another.

Reasons are keyed to methodology. Among the factors that figure in whether others can responsibly reject the claim that c supports p , are views about methods of establishing whether p . Are standardized aptitude tests a good predictor of college success? Or can one reasonably reject the contention that a student’s acing the SATs indicates that he will do well? Are focus groups a good way of gauging public opinion? Or can one reasonably reject the contention that because a focus group found the candidate’s message compelling, the public at large will do so too? To answer such questions we need to know whether the methods that connect c and p have been validated, and what standards of rigor apply in cases of a particular kind. That is, we need to know what is known or reasonably believed about a topic and the methods for

investigating it.

A major question is whose reactions matter? Who is comprehended in the ‘we’? To be exceedingly tolerant, to allow just about everyone to have a say, is almost surely to court skepticism. For any interesting thesis, there are bound to be people whose background beliefs are such that they could, without violating their own epistemic principles, reject a contention that c is a reason to believe that p . A blind person might reject the contention that the car’s looking blue in daylight is reason to believe that it is blue. He has no direct evidence of that. A novice, seeing that the university admits more men than women, might reject the contention that a university does not discriminate against women, because he does not understand the effect on statistics of properly partitioning a domain (Cartwright, 1983, pp. 36-38). He does not understand why, if departments do their own admissions, looking at the department level rather than at the whole university yields the appropriate statistics. A religious fundamentalist might reject the contention that the fossil record affords reason to believe that birds evolved from dinosaurs. He denies that anything could afford genuine reason to believe that evolution occurred. To disqualify such reactions requires a principled way to delineate the class of others whose reactions matter.

More is needed than that the standards be shared. All three rejecters might be deploying standards they share with others. The issue is whether those standards are cognitively good standards, standards whose satisfaction fosters the growth of knowledge or the advancement of understanding. To resolve this, we need to look at the standards and the goods they promote.

We dismiss the rejecters’ opinions because we know better. This is not

intellectual arrogance. We know what relevant resources are available and how and why to draw on them. We dismiss the statistical novice's opinion because it rests on a misunderstanding of statistical methods. This indicates that reasons must consonant with the proper use of acceptable methods, where the acceptability of the methods turns on their having been validated and endorsed by those who understand the subject and the best available ways to justify conclusions about it. But it is not enough if, by chance, statistics are properly used. The proper use of the methods should be based on an understanding of the methods, their ranges of application, their powers and limitations. We dismiss the blind man's opinion because it is based on a dearth of direct evidence and a failure to make proper use of available indirect evidence. This indicates that we are required to be aware of our own limitations and of the resources we can draw on to compensate for them. We dismiss the fundamentalist's stance because it is dogmatic. It is not sensitive to evidence, since the fundamentalist would retain his position whatever the evidence. This indicates that for something to count as a reason, it should be responsive to evidence.

The fundamentalist might reply that his opinion is responsive to evidence. Scripture, he maintains, provides incontrovertible evidence that evolution did not occur. He disregards the fossil record because nothing it shows could override scripture. His belief then rests on shaky grounds. He relies on a single source, so his belief is unstable. He relies on a single, disputed interpretation of the deliverances of that source, so it is frail.

The scientist is in an epistemically stronger position. Her belief is solidly grounded. We would have to be massively wrong about genetics, anatomy, and

physiology to discredit the direct biological evidence for evolution; about physics, to discredit the evidence provided by carbon dating; and about geology to discredit the evidence of the fossil record. The scientist then can draw on a sophisticated theory of empirical evidence, which explains why and how the sorts of factors biology relies on should be considered trustworthy. A second strength to her position is its fallibilism. Science does not consider any methods or results incontrovertible. Even the best of today's methods and findings might need to be reconsidered on the basis of future findings. So science builds in solidifiers to accommodate its recognition that with the advancement of understanding come refinements in methods and epistemic standards. The conviction that standards and methods neither can nor need ever be revised leaves the dogmatist in an epistemically precarious position.

It may seem paradoxical that beliefs we are prepared to revise are epistemically more solidly grounded than ones its adherents consider unrevisable. But that is so. The epistemic value of science lies not in the conclusiveness of its results, but in the self-correcting character of its methods and standards. Because we consider a method good, we consider its deliverances trustworthy. Because it delivers results that satisfy our epistemic standards, we consider it a good method. Because results that satisfy these standards promote our epistemic ends, we consider these standards the proper ones. Our views about the particular theses we take ourselves to have reasons for intertwine with the methods we take to count as generating reasons, and the standards we take reasonable beliefs to have to satisfy. As Walden (2007) says, 'We never properly get evidence for a proposition; we get evidence for a proposition relative to methodological prescriptions about what counts as evidence for what. . . . These methodological prescriptions are also

part of our theory, and thus . . . *they* are actively confirmed or denied alongside the rest of the theory.’ The interdependence is dynamic. As our understanding of a topic grows, so does our understanding of the best ways to investigate it, to validate the methods for investigating it, to identify the standards to which views about it should be held, and hence to characterize what counts as evidence in that area. Solidity does more than just provide insurance. It also deepens understanding. The various strands tie a solid belief to different other commitments, yielding a wider perspective as to where it fits and how it functions in our system of beliefs.

In *Truth and Truthfulness*, Williams shows that truthfulness is a thick concept whose extension depends in complicated ways on the value of conveying truths to one another. I have given an almost equally complicated account of trustworthiness. Our cognitive goal, I have urged, is not to ace life’s true/false test. It involves forming beliefs (and other attitudes) that we can use as a reasonable basis for inference and action and can responsibly convey to others when interests are cognitive. My justified true beliefs or reliable true beliefs fail to be trustworthy if their grounds are ones that other members of the epistemic community can reasonably reject. In that case, the grounds do not constitute reasons. But if I am a member of an epistemic community, I share its cognitive values. So I should not on reflection endorse conclusions that my similarly situated fellows would reasonably reject. Although grounds are objective, reasons and trustworthiness are intersubjective. Community standards determine how solidly grounded my beliefs should be and what supplies the grounds. The answers to such questions change with the advancement of understanding in a field, the development of new methods for investigating it, and the revision of relevant standards of acceptability.

I have not argued that trustworthiness is integral to knowledge, justified belief, or reliable belief. Rather, I have urged that it figures in what makes knowledge, justified belief, and reliable belief worth having. Its thickness derives from our inability to delimit the class of trustworthy beliefs without considering what we want true beliefs for, what makes particular true beliefs worth having.

References

Bonjour, Laurence. (1985). *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.

Cartwright, Nancy. (1983). 'Causal Laws and Effective Strategies,' *How the Laws of Physics Lie*. Oxford: Clarendon Press.

Scanlon, T. M. (1998), *What We Owe to Each Other*. Cambridge: Harvard University Press.

Walden, Kenneth (2007), 'A Practico Priori,' unpublished manuscript.

Williams, Bernard. (2002). *Truth and Truthfulness*. Princeton: Princeton University Press.