

Is Fallibility an Epistemological Shortcoming?

External world skepticism appears to be plainly false. Even in our most careful everyday investigations of whether people know particular things we often arrive at the conclusion that they undoubtedly do. So if we are seriously to entertain the possibility that no one can know anything about the world, we need to be given an argument in its favor. The goal of this paper is to refute one familiar form of skepticism in order to deepen our understanding of what such an argument would have to be like.

It is often said that skepticism's source is the idea that knowledge requires infallibly true belief. For instance, Dretske, Lewis, and others have held that skepticism turns upon what I will call the *Infallibility Requirement*, the requirement that in order to know something about the world, one must be able to "rule out" or "eliminate" every possible way in which one could be wrong.¹ "Rule out" and "eliminate" have been interpreted in various ways. Any interpretation will do for my purposes, so long as satisfying the Infallibility Requirement yields infallibly true belief. For this reason, it is important not to equate "ruling out" with "knowing (or being in a position to know) not to obtain,"² since on this interpretation meeting the requirement yields infallibly true belief only if we assume that knowledge requires infallibly true belief. Likewise, one should not equate the Infallibility Requirement with a closure principle for knowledge.

We seldom, if ever, meet the Infallibility Requirement. No matter how good our evidence, it always leaves open ways in which our beliefs about the world could be wrong.

¹ David Lewis, "Elusive Knowledge," *Australasian Journal of Philosophy* 74 (1996), pp. 549 – 67, at p. 549; Fred Dretske, "The Pragmatic Dimension of Knowledge," *Philosophical Studies* 40 (1981), pp. 363-78, at p. 365.

² As Dretske, for instance, does (*ibid.*, p. 331).

So if the Infallibility Requirement is correct, our fallibility provides a decisive argument for skepticism. In what follows, I will call this skeptical view "Infallibilist skepticism," its proponent the "Infallibilist skeptic."³

The Infallibility Requirement plays no role in our ordinary practices of knowledge attribution, even when we are being conscientious and careful. This fact provides a simple objection to Infallibilist skepticism. Since we ordinarily say that people know things even when they don't meet the Infallibility Requirement, hasn't the Infallibilist skeptic misunderstood the requirements that we must meet in order to have knowledge? As I will argue, this objection is correct.

This objection is reminiscent of ordinary language philosophy, particularly J. L. Austin.⁴ According to one common interpretation of that tradition, its guiding idea was that our ordinary linguistic practices directly determine the requirements which must be met in order for the ascription of any given predicate to be true. On such a view, skepticism would be directly refuted by the fact that even when we are being careful and conscientious we often say that people know things. I do not accept this view. Consequently, my argument will not presuppose or attempt to establish that our ordinary linguistic usage directly determines the conditions for knowledge. Rather, I will consider whether our ordinary practices of epistemic evaluation — the requirements which we ordinarily deploy and the conditions under which we ordinarily think it correct to ascribe knowledge to ourselves and others — are a *good guide* to the requirements which we must meet in order to possess knowledge. My aim is to show that it is reasonable to think they are, and that

³ For a defense of Infallibilist skepticism, see Peter Unger *Ignorance* (New York, NY: Oxford University Press, 1975). (Unger has since changed his mind.)

⁴ "Other Minds," reprinted in his *Philosophical Papers* (J. O. Urmson and G. J. Warnock, eds., Oxford: Oxford University Press, 1961, pp. 44 - 84), and *Sense and Sensibilia* (G. J. Warnock, ed., Oxford: Oxford Univ. Press, 1962).

we therefore have good reason to reject any skeptical argument which deploys requirements which are not found in our ordinary epistemic practice. Thus while my immediate target is Infallibilist skepticism, its failure enables a more general conclusion. A convincing skeptical argument must use only requirements to which we are committed by our ordinary practices.

1. Infallibilism's Explanatory Task

Any external world skeptic owes us an explanation. If no one can ever know anything about the world, why do we confidently say and believe that we know things? To see the significance of this question, consider the confrontation between the Infallibilist skeptic and our ordinary knowledge attributions.

Despite our inability to meet the Infallibility Requirement, we ordinarily say and believe that people possess knowledge. Even in our most careful everyday investigations of whether people know things we would find an Infallibilist denial of these knowledge attributions to be outrageous. It is part of our ordinary practices that we take these responses to be backed by good reasons; we often take ourselves to have excellent reasons for attributing knowledge to people and at least sometimes think that there is no reason to doubt these knowledge attributions. So insofar as the Infallibilist skeptic purports to be talking about the epistemic status which is at issue in our ordinary epistemic evaluations, there is, *prima facie*, good reason to doubt the correctness of the Infallibility Requirement. Since we do not ordinarily insist upon this requirement, it is reasonable to suspect that the Infallibilist is either changing the subject or misinterpreting some ordinary requirement. If he is to allay these suspicions, the Infallibilist must provide us with some reason not to confidently acquiesce in our ordinary epistemic judgments. Insisting upon the correctness of the Infallibility Requirement is not enough.

There is a standard way of undercutting this sort of anti-skeptical appeal to our ordinary practices of knowledge attribution. This is to draw a conceptual distinction

between the conditions under which it is appropriate to call certain cases "cases of knowledge", on the one hand, and the conditions which people must meet in order to actually know things about the world, on the other.⁵ This move is perfectly correct, so far as it goes. There is, in general, a conceptual distinction between saying something that is appropriate or reasonable in one's circumstances and saying something that is true. It is often appropriate or reasonable to say something that is false (for instance, if one has good evidence for it and is ignorant of its falsehood), and it is sometimes inappropriate or unreasonable to say something that is true. Consequently, the mere fact that we ordinarily say, quite reasonably and appropriately, that people know things does not entail that they really do; the conditions for appropriately claiming knowledge and for actually having it may differ.

It has been maintained that to overcome skepticism by appealing to our ordinary practices, one must deny this distinction.⁶ This is incorrect. The fact that knowledge attributions which are appropriate are not *thereby* true is perfectly compatible with the possibility that many of our appropriate attributions of knowledge *are* true, and it does not provide any reason to think that they are false. Since we ordinarily take ourselves to have

⁵This response is inspired by H. P. Grice's arguments against the methods (as he understood them) of ordinary language philosophy. See, in particular, "Logic and Conversation" and "The Causal Theory of Perception," both in *Studies in the Way of Words* (Cambridge, MA: Harvard University Press, 1989). Peter Unger, *Ignorance*, pp. 50–54, and Barry Stroud, *The Significance of Philosophical Scepticism* (New York, NY: Oxford University Press, 1984), chapter 2, offer a broadly Gricean defense of skepticism, discussed below.

⁶ Stroud, *ibid.*, p.64. Stroud maintains, moreover, that one is consequently forced to deny the platitude that the world is as it is regardless of how we think, believe, or say it is and regardless of whether or not we can know how it is (*ibid.*, pp. 76 ff.).

very good reasons for concluding that people possess knowledge about the world and no reason to believe that they don't, we appropriately conclude that people know things about the world. But, of course, they don't meet the Infallibility Requirement. So we may reasonably conclude that the Infallibility Requirement is incorrect. We can thus simultaneously countenance the conceptual distinction between appropriate and true knowledge attributions and also reject the Infallibility Requirement on the basis of our ordinary epistemic practices. So if the Infallibilist skeptic is to move us from our position of ordinary confidence that we have knowledge, he must do more than merely invoke this distinction. He must provide us with some reason to think that the considerations which guide our ordinary knowledge attributions do not fully or correctly reflect the requirements for actually possessing knowledge.

Such a reason will not be provided if the Infallibilist simply claims special insight, denied the rest of us, into the requirements dictated by the concept of knowledge itself quite apart from our practices. For two can play this game. Why can't we reply simply that *our* insight into the concept of knowledge reveals no such requirement, or that the Infallibilist has mistakenly latched onto the wrong concept? In order to make his case, the Infallibilist must appeal to considerations about our ordinary practice itself. In particular, he must *explain it away*. For why, if the Infallibility Requirement is correct, is it nonetheless appropriate for us to ascribe knowledge to people even though we recognize that they can't meet it? An answer adequate to the Infallibilist's purposes would show that our ordinary knowledge ascriptions are responsive to considerations which have nothing to do with their truth. Such an explanation would show that our practice does not fully or correctly reflect the requirements for knowledge and would thus give the Infallibilist a strong case for dismissing the appeal to our ordinary practices. And we should demand such an explanation anyway. Any view which maintains the falsity of a range of assertions which we treat as perfectly appropriate should be prepared to explain why we treat those assertions as appropriate despite their falsehood.

To discharge this explanatory burden, the Infallibilist must explain why it is reasonable or appropriate for us to *wave* or *ignore* the Infallibility Requirement in the course of our ordinary procedures of knowledge evaluation. The best attempt in this general direction has been made by Barry Stroud, following a suggestion of Peter Unger's. Stroud proposes that although the concept of knowledge — the concept that guides our everyday epistemic assessments — involves certain unmeetable requirements, we ignore this fact because of the practical and social circumstances in which we ordinarily make and assess knowledge claims. We do so, Stroud suggests, because meeting the requirements that we ordinarily impose puts us in a position that is close enough, for all practical purposes, to knowledge.⁷

Stroud's suggestion is motivated by the thought that our ordinary epistemic activities are tied up with our practical concerns. Our pursuit of the truth in everyday life is constrained by practical interests and circumstances; we are hampered by limitations of time and resources, and on some occasions the truth matters more to us than on others. Gathering evidence in order to eliminate competing possibilities, asserting that one knows something, and the like are all activities or actions. So they are susceptible to practical evaluation — evaluation in terms of the reasonableness of doing them (rather than something else) given one's practical situation, one's purpose, the time one has available, etc. Thus, even if one lacks knowledge, it may be reasonable, in practical terms, to cease gathering evidence and to claim knowledge.⁸ Stroud suggests, therefore, that the reasonableness or appropriateness of our ordinary knowledge claims can be understood as

⁷ Stroud, *ibid.*, pp. 64 ff., especially pp. 71 –2; Unger, *Ignorance*, pp. 50-54. Stroud is not concerned to defend Infallibilist skepticism in particular; his suggestion, if successful, would also defend other forms of skepticism against the objection from our ordinary practice.

⁸For Stroud's statement of the argument, see *op. cit.*, p. 66.

practical reasonableness; because of the necessities of our practical lives, we waive certain requirements when applying the concept of knowledge. I will hereafter call this account the *Practical Constraints View*, since it holds that our application of the requirements dictated by the concept of knowledge is constrained by the practical circumstances of our ordinary knowledge evaluations.

If this view were true, it would provide the skeptic with a tidy explanation of how we discover the truth of skepticism. For this view encourages us to understand philosophical reflection as a matter of stepping back from our practical concerns in order to gain a clear view of the conditions for the true application of our concepts. According to the Practical Constraints View, then, the truth of skepticism is revealed when we reflect on our epistemic concepts in isolation from the practical constraints governing their ordinary applications.⁹

The Practical Constraints View offers a coherent and attractive vision of the relation between skepticism and our ordinary knowledge attributions. However, we also need some reason to think that it is true. Otherwise, the Infallibilist's position would collapse; his "explanation" of our practice wouldn't be any explanation at all. Since the Practical Constraints View is a theory about our ordinary practices of knowledge attribution (in particular, about the conditions under which knowledge claims and attributions are taken to be appropriate), we can test its plausibility by investigating whether it offers an accurate description of our actual practices. If we find it to be incorrect, then we will have no reason to take Infallibilist skepticism seriously. We will be free to reject Infallibilism by appealing to our ordinary practices.

⁹cf. Stroud, *op. cit.*, pp. 71 ff. Stroud draws upon Thompson Clarke's "The Legacy of Skepticism" (*The Journal of Philosophy*, 69, 1972, pp. 754 – 69), in which Clarke discusses this conception of philosophical reflection at length, though in the end he doubts its full intelligibility.

2. *The Failure of the Practical Constraints View*

Stroud defends the Practical Constraints View by means of an example. As I will now argue, however, his interpretation of this example is incorrect. To put the point roughly, even in the press of practical circumstances we do not think it appropriate to waive or ignore requirements for knowledge or to claim knowledge if one's epistemic position is merely adequate for practical purposes. Instead, we say (in effect), "Knowledge, shmowledge! We need the best judgment available and have to be content with that."

Stroud's example is as follows.¹⁰ Imagine a group of soldiers who have been trained to visually identify enemy aircraft from the ground. They have been taught, and their training manual states, that aircraft exhibiting features x, y, and z are of type F. However, there are also enemy aircraft of another type, type G, which are indistinguishable from F's when observed from the ground. The airplane spotters were not taught about G's because the existence of G's is irrelevant to the war effort; they are rare, antiquated, and harmless, while F's are extremely dangerous. Telling the spotters about G's would present needless complications, both during their training and in the field.

Stroud makes three correct observations about this example. First, if a spotter determines that a plane flying overhead has features x, y, and z, then it is appropriate or reasonable for him to claim to know that it is an F. Second, such a spotter does not in fact know that the plane is an F. For all he knows, it might be a G. Third, there is no good reason, in the context of the war effort, to tell the spotter that he doesn't know that the plane is an F; doing so would have no practical point.

Stroud also links these three observations, proposing that since there is no practical point in telling the spotters that they lack knowledge, the requirement that they eliminate the

¹⁰ *Op. cit.*, pp. 67 ff. The example is adapted from an example of Clarke's, *op. cit.*, pp. 759 ff.

possibility that the plane is a G has been waived *for practical reasons*. He thus takes the spotters' practical circumstances to explain *both* why it is pointless to challenge a spotter's knowledge claim *and* why it is appropriate for the spotter to claim, falsely, to know that a plane is an F. His thinking here seems to be as follows. Given the wartime setting, the dangerousness of planes of type F makes it very important that a spotter be right when he or she claims that a plane is not an F. But it is not so important that a spotter be right when he or she claims that a plane *is* an F, since it is better to shoot down an occasional harmless plane than to let a dangerous F get through. So given the rarity of G's and the dangerousness of F's, practical considerations warrant ignoring the possibility that a plane is a G when trying to determine whether it should be shot down as an F; for all practical purposes, the possibility that a plane is a G is irrelevant and may be ignored. Consequently, a spotter may appropriately claim to know that a plane is an F even if he has not eliminated or even considered the possibility that it is a G.

If this is the correct interpretation of the example, then it would be plausible to hold that practical considerations play a similar role in our practices of knowledge attribution more generally. Consequently, we could not infer the incorrectness of the Infallibility Requirement from the fact that it makes no appearance in our ordinary practices.

But is this the correct interpretation? If, as Stroud urges, the spotters' knowledge claims are appropriate or reasonable because the existence of G's is irrelevant in the practical circumstances, then it should also be appropriate for them to ignore the possibility that a given plane is a G and claim knowledge that it is an F *even if* they know about the existence of G's. However, this implication is incorrect. If one knows about the existence of G's, it is not appropriate, even within the context of the war effort, to ignore the possibility that a plane is a G when one claims to know it is an F or attributes such knowledge to someone else. To see this, imagine, first, that you are someone (say, a general) in the spotters' context who knows about the existence of G's. Given the practical context, you would quite appropriately act on the information provided by a conscientious

spotter. But would you feel that it is appropriate to say that the spotter *knows* the plane is an F? I doubt you would, or at least if you did, you would also feel that you should be prepared to qualify and explain the remark. It would not be appropriate simply to *ignore* the very real possibility that a plane is a G when making claims about what the spotters know. Likewise, imagine that you are a spotter who has been told about the existence of G's. You would not feel that it is appropriate to claim knowledge that a plane is an F, even in the thick of battle. You might have no qualms about declaring a plane to be an F. But you wouldn't claim to know it, or at least if you did, it would be with a sense that what you are saying isn't unobjectionable as it stands but requires qualification and explanation (which you may not have time to give). Consequently, the appropriateness of the spotters' knowledge claims isn't explained simply by the fact that the existence of G's is irrelevant in their practical circumstances.

In fact, practical considerations have nothing to do with the reasonableness of the spotters' belief that they have knowledge. To see this, consider a slightly different example.¹¹ An 18th-century ornithologist is attempting to catalogue the species of birds present in a certain area. According to the classificatory standards accepted at the time, a bird which exhibits characteristics a, b, and c while in flight is of species M. However, there is a very rare species of birds, N, which has not yet been identified and is unique to the area. Birds of this species also exhibit features a, b, and c while in flight and are otherwise indistinguishable (from the ground) from birds of species M. Upon observing that a bird in flight exhibits features a, b, and c, the ornithologist might claim to know that there is an M in the area, and both we who know about N's and his colleagues (who don't) would regard his knowledge claim as being perfectly reasonable and appropriate. Still, he does not know that the bird is an M. For all he knows, it is an N.

¹¹I am indebted here to discussion with an undergraduate class at Harvard University, and particularly to Paul Monteleoni.

What makes it reasonable or appropriate for the ornithologist to claim to know that the bird is an M? It seems that the ornithologist's knowledge claim is reasonable just because he has good reason to conclude that he knows the bird to be an M. Like all ornithologists at the time, and through no fault of his own, he was simply ignorant of the existence of N's, and so he failed to realize that features a, b, and c were inadequate for establishing that a bird is an M. But given the state of his knowledge, he proceeded impeccably on the basis of the evidence available to him. Being aware that this was so, he quite reasonably claimed to know that the bird was an M. Analogous remarks can be made to explain why it would often be reasonable for the spotters, ignorant of the existence of G's, to claim, falsely, to know that a certain plane is an F. Their knowledge claims would be reasonable or appropriate because (1) they are ignorant through no fault of their own of the existence of G's, (2) they are consequently aware of no reason to think that their evidence is inadequate, (3) they proceeded impeccably on the basis of the information available to them, and (4) they are aware of having done so. None of this is a practical matter.

It is important at this point to distinguish the practical or conversational factors which govern what it would be sensible to *say* in a certain setting and the evidential considerations which govern what it would be epistemically appropriate or reasonable to judge, conclude, or believe. Practical and conversational considerations obviously affect what it is sensible to say in the airplane spotters' context. The reason that there is any point in a spotter's declaring "I know that it is an F" is that he is engaged in the activity of attempting to identify planes as they fly overhead. Likewise, it would not serve the purposes of the war effort to explain the existence of G's to the spotters or deny their claims to know that certain planes are F's. So Stroud is right that when we judge that it would be unreasonable or inappropriate to point out that the spotters lack knowledge, what we have in mind is practical or conversational inappropriateness. However, the fact that it would be conversationally or practically inappropriate to point out their lack of knowledge

does not show that it is *for practical reasons* that the spotters are reasonable in taking themselves to have knowledge. When we judge that it is reasonable for them to claim knowledge, what we have in mind is, in the first instance, *evidentiary* or *epistemic* appropriateness. Like the ornithologist, they are epistemically reasonable in concluding or believing that they know. It is for this reason that it is appropriate or reasonable for a spotter to assert that he knows the plane is an F when he is in a setting in which this remark would be germane.

Correctly interpreted, then, the example of the airplane spotters simply reminds us that if one is unaware of certain facts through no fault of one's own, then one can sometimes be epistemically warranted in claiming to know something even though one actually doesn't know it. This lesson is simply an instance of the general principle that one can be epistemically warranted in believing that the conditions for the truth of an assertion are met and yet still be wrong — the principle applies to assertions that one knows something as much as to any other. Consequently, the example fails to establish Stroud's claim that it is sometimes appropriate or reasonable for purely practical reasons to ignore certain possibilities when one makes a knowledge claim. It is not for practical reasons that having good evidence for the truth of p can entitle one to conclude that p even when p is false.

In fact, the example of the airplane spotters supports a conclusion directly opposed to the Practical Constraints View. Consider that regardless of the practical circumstances, a spotter who knows about the existence of G's would not think it appropriate or reasonable to claim knowledge that a certain plane is an F. This is an example of a widespread phenomenon. When a fact is pointed out to us which is admittedly irrelevant for practical purposes but relevant to the question of the truth of what we claim to know, we do not simply ignore it and continue to claim knowledge. Instead, we think it unreasonable or inappropriate to continue to claim knowledge unless we can do something to show that the alternative in question does not obtain. This strongly suggests that our everyday epistemic

evaluations attempt to track the conditions for the *truth*, not the practical appropriateness, of our knowledge claims. When we claim knowledge, we — like the spotters — do so because we think we are warranted in concluding that the conditions for the truth of the knowledge claim are met. The reasonableness of our everyday knowledge attributions is thus primarily epistemic, not merely practical: we're trying to say what is both true and practically or conversationally pertinent, not what is merely appropriate or useful for practical purposes. Of course, people sometimes *say* that they know something even when they don't believe that they do. Mendacity in the service of practical goals is at least as common here as elsewhere. But this fact is not relevant to our discussion. Our question concerns the considerations which guide our sincere and conscientious evaluations of the state of people's knowledge. As the example of the spotters reminds us, practical considerations don't affect the requirements which we insist upon in the course of such assessments.

3. Infallibilism, Practical Considerations, and Our Epistemic Ideals

I now want to deepen and defend this conclusion by showing in detail that Infallibilist skepticism founders upon an incorrect conception of the relation between practical considerations and the standards for appropriate knowledge attribution. This discussion will lead to a more plausible account of the considerations which guide our ordinary epistemic practice and will provide us with strong positive reasons for rejecting Infallibilist skepticism.

Consider the following account of our attributive practice:

The standards which a person must meet in order for it to be appropriate to attribute knowledge to him or her vary with the practical and conversational context. In ordinary circumstances, we do not require people to be able to eliminate all possibilities of error. Instead, we deem it appropriate to attribute knowledge to them if they merely attain a position close enough for

current practical and conversational purposes to being able to eliminate all possibilities of error, for instance by being able to eliminate a good many of them, especially the relevant or salient ones. However, in some contexts the standards are considerably more stringent, and in some contexts (such as the context created by philosophical reflection) a person must be able to rule out all possibilities of error in order for it to be appropriate to attribute knowledge to him or her.

This is the account of our attributive practice which results when the Infallibilist explicitly adopts the Practical Constraints view. Infallibility, on this view, plays a crucial role in our practices of epistemic evaluation; it is the ideal against which we measure a person when we consider whether to attribute knowledge to him or her. But despite its emphasis upon Infallibility as our epistemic ideal, this account is not distinctive to Infallibilist skepticism. It is shared by the "contextualist" accounts of the semantics of knowledge attributions which have recently gained prominence as a reply to Infallibilist skepticism, and its detailed development is found primarily in these responses.¹² The contextualist's basic response to Infallibilism is to accept the above account of the *appropriateness* conditions for knowledge attribution, and then to claim that these conditions are also the conditions for the *truth* of

¹²See, in particular, David Lewis, "Elusive Knowledge"; Keith DeRose, "Contextualism and Knowledge Attributions" (*Philosophy and Phenomenological Research* 41 (1992), pp. 913-29), "Solving the Skeptical Problem" (*The Philosophical Review*, vol. 104 (1995), pp. 1-49), and "Contextualism: An Explanation and Defense" (in *The Blackwell Guide to Epistemology*, J. Greco and E. Sosa (eds.), Malden, MA: Blackwell Publishers, 1999, pp. 187 – 205); and Stewart Cohen, "How to be a Fallibilist" (*Philosophical Perspectives*, 2, Tomberlin (ed.), 1988, pp. 581 –60), "Contextualism, Skepticism, and the Structure of Reasons" (*Philosophical Perspectives*, 13, Tomberlin (ed.), 1999, pp. 57 – 89), and "Contextualism and Skepticism" (*Philosophical Issues*, 10, 2000, pp. 94 – 107).

knowledge attributions. On the resulting view, the standards for possessing knowledge will shift with the context, and since they are low in ordinary settings, what one says when one claims or attributes knowledge in such settings may well be true. From here on, I will use "Practical Constraints Contextualism" as a label for the account of our attributive practice shared by the Infallibilist and his contextualist opponent. (This account only concerns the appropriateness conditions for knowledge attributions. It does *not* involve the further semantic claim that the truth conditions of knowledge attributions shift with the practical and conversational context.) Practical Constraints Contextualism is badly mistaken, as I will now argue.

In what follows, I will use the term "relevant error possibilities" to refer to the possibilities of error which a person must be able to eliminate in order for it to be appropriate to attribute knowledge of a given proposition p to him or her in a given context. By "possibility of error," I mean any possibility which is incompatible with p or with the person's knowing that p .

Two things are needed if we are to have good reason to accept Practical Constraints Contextualism. First, we need some plausible examples of shifts in the set of relevant error possibilities. Second, we need an account of the mechanism(s) by which these shifts are effected.

Examples of such shifts do not seem hard to come by. For instance, in ordinary circumstances in which someone claims to have seen a goldfinch, it would be thoroughly inappropriate to object, "But you don't know whether that is a bird at all; it might just be a very clever hologram." In other circumstances in which the person's evidence is exactly the same, this response would be quite appropriate; we would not attribute knowledge to the person unless she had extensive and specific evidence against this possibility. So the conditions for appropriate knowledge attribution do appear to shift. The crucial question is whether these shifts are due to practical and conversational factors, as practical constraints

contextualism asserts. In order to answer that question, we need a detailed proposal about how such factors might govern these shifts.

According to one common proposal, conversational salience is the primary source of these shifts: the mere mention of a possibility of error tends to make it conversationally relevant and thus to raise the standards for appropriate knowledge attribution.¹³ This claim is incorrect. Suppose that while on an ordinary walk in the woods, you claim to see a goldfinch. Your friend suggests — without any reason — that it might just be a clever hologram. You will not respond by saying (or thinking), "Now that he's mentioned this possibility I shouldn't claim to know it's a goldfinch, since I don't have any specific evidence that it's not just a hologram." Nor will you feel obliged to investigate the matter (for instance, by rummaging through the underbrush in search of an apparatus and a source of electrical current). Instead, you will respond, "Don't be silly!" If it was appropriate for you to claim knowledge before the suggestion was made, it is appropriate afterwards as well. Mere mention of an outrageous possibility does not change the standards at all. This is not to say that it would be appropriate for you to assert such things as, "I know that it is a goldfinch, but I don't know that it is not just a clever hologram." It would be appropriate for you to claim to know it's not a hologram, as well.

¹³ See in particular Lewis' "Rule of Attention:" "No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in *this* context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative" ("Elusive Knowledge," p. 559). Cf. DeRose, "Solving the Skeptical Problem," p. 36 fn. 34; Cohen, "How to be a Fallibilist," p. 96. DeRose ("Solving," *passim*) offers a minor variant, suggesting that the shift is induced mainly by mention of the possibility within the scope of an epistemic operator, as in "You don't know that that's not just a hologram." This doesn't change the fundamental issue.

It is sometimes suggested that one can “resist” the rise in standards which is putatively induced by the mention of an error possibility.¹⁴ However, it would be incorrect to appeal to resistance to account for the example I just described. Talk of “resistance” requires that a conversationally-induced shift in the standards would at least be unexceptionable in this case. But if conditions are normal, it would be irrational and bizarre to worry about the possibility of holograms, even after that possibility has been explicitly mentioned. Just imagine an experienced birdwatcher who, while walking through the woods and without any reason to suspect deception, refused to say that he knew a certain bird was a goldfinch until he had thoroughly investigated the setting in order to insure that what he had seen was not a hologram. Such behavior would be quite odd, to say the least. To find a plausible version of Practical Constraints Contextualism we must look elsewhere.

According to the most promising current suggestion, the set of relevant error possibilities for a given proposition is determined by or relative to practical features of the context of attribution – such features as the purposes and interests of the evaluators, the

¹⁴ "Elusive Knowledge," p. 560. DeRose suggests something similar when he notes that not every mention of a skeptical hypothesis will succeed in raising the standards for knowledge ("Solving the Skeptical Problem," p. 15 fn. 22, p. 36 fn. 34.)

cost of error, the practical limitations and necessities which are operative, etc.¹⁵ For instance, it is often suggested that in a context in which a great deal is at stake, such as a courtroom, the standards for appropriate knowledge attribution will shift to include possibilities of error which would be ignored in more quotidian settings. The basic idea here is this. In each particular context, the total error possibilities for a given proposition are ordered such that certain alternatives are in some sense "closer" or "more relevant" than others. The interests and practical limitations of the relevant people in the particular context then set a context-specific standard for appropriate knowledge attribution by selecting a *range* of possibilities that someone must be able to rule out.¹⁶ The more it matters that *p* be true (whether because of the relevant people's practical interests or the purity of their desire to determine the truth of *p*), the wider the range of alternatives the subject of appraisal must

¹⁵ Lewis, "Elusive Knowledge," p. 556; Unger, *Philosophical Relativity* (Oxford, England : Basil Blackwell, 1984, at p. 48); DeRose, "Contextualism: An Explanation and Defense," p. 191, "Solving the Skeptical Problem," p. 10 fn. 14, and the example discussed in "Contextualism and Knowledge Attributions," pp. 913 ff.; Cohen, "Contextualism, Skepticism, and Reasons," p. 61 (and the example discussed on pp. 58-9); Fogelin, *Pyrrhonian Reflections on Knowledge and Justification* (New York, NY: Oxford University Press, 1994), at p. 198; Stine, "Skepticism, Relevant Alternatives, and Epistemic Closure" (*Philosophical Studies*, 29, 1976, pp. 249 – 61), at p. 254.

¹⁶ The idea of an ordering of alternative possibilities (and the conception of "epistemic positions" as being determined by the range of alternative which one can rule out) is explicit in DeRose's gloss of his "Rule of Sensitivity" in a possible worlds framework ("Solving the Skeptical Problem," p. 37). A similar conception is also implicit in Lewis' discussion. The intuitive idea of a relevance-ordering could be given many different theoretical interpretations. Furthermore, the ordering may plausibly be treated as only a partial ordering and need not even be fully determinate.

be able to rule out in order for it to be appropriate to attribute knowledge to him or her. Other factors might also be held to make other specific error possibilities salient.¹⁷ But this broad mechanism would account quite generally for contextual shifts in the set of relevant error possibilities and thus would provide a plausible way of cashing out the Infallibilist's talk of being "close enough for practical purposes" to being able to rule out every possibility of error.

Is the idea of an ordering or ranking amongst alternatives *essential* to this version of Practical Constraints Contextualism? It might be urged that the practical context simply determines what (rough) proportion of the total set of error possibilities one must be able to rule out. However, in a given context, certain quite particular error possibilities — but not others — will be relevant. In an ordinary situation it would do you no good to rule out the possibility that a putative goldfinch is an intergalactic spying device if you couldn't even show that it wasn't just a bird of some other common, similar-looking type. The suggestion that the context simply selects a *proportion* of the total alternative possibilities fails to explain this. As I argued above, however, the appeal to conversational salience is also inadequate. So I don't see any way, short of assuming an ordering of possibilities and a contextually set range within that ordering, for the Infallibilist skeptic to offer a plausible position. I will consequently assume that a relevancy-ordering of alternatives is integral to any plausible version of Practical Constraints Contextualism.

The proposed view involves two crucial commitments. The first is that the range of error possibilities that a person must be able to eliminate will vary with the interests of the relevant people and the nature of the practical context. The second is that when we assess someone's putative knowledge that *p*, we are concerned to determine the range of error

¹⁷ See, for example, the rules Lewis proposes to govern contextual relevance in "Elusive Knowledge." (Of course, I have already rejected his "Rule of Attention".)

possibilities that the person is able to rule out. As I will now argue, both ideas are incorrect.

In order to make my case, I will describe some examples and invite you to consider whether, were you in the envisaged circumstances, you would judge a particular knowledge claim to be appropriate. This procedure is legitimized by the fact that Practical Constraints Contextualism is itself a theory of the conditions under which it is appropriate to attribute knowledge. Since these conditions are purely a matter of our attributive practices, the theory must attempt to capture the conditions under which we would judge that a knowledge attribution was appropriately made. Accordingly, it can be tested against our actual judgments. It might be objected that this will not be probative, since my opponent may have opposing intuitions.¹⁸ However, since our judgments are an exemplification of our practices of epistemic evaluation and knowledge attribution, what matters here are our actual responses to the examples. Admittedly, a philosopher who has a theoretical commitment to Practical Constraints Contextualism will have to claim opposing judgments. However, someone who embodies Practical Constraints Contextualism, who lives his or her epistemic life in accordance with it, proceeds in an extremely strange manner, as my examples are meant to suggest.

I do not take my examples to be decisive counterexamples. My aim is rather to urge that in its most basic form, Practical Constraints Contextualism fundamentally misrepresents our ordinary procedures of knowledge attribution. It could perhaps be made to fit the data, but its basic idea is on the wrong track.

a. Practical Conditions and Standards for Appropriate Knowledge Attribution

We have already seen an example which strongly suggests that practical conditions do not *constrain* the range of relevant error possibilities. Recall the case of the airplane

¹⁸ I am grateful to Steven Gross and Jim Pryor for raising this issue.

spotters. In their situation, the possibility that a particular plane is a G is irrelevant for practical purposes, and it is not practically feasible to attempt to rule it out. Hence, according to Practical Constraints Contextualism it should be appropriate for an airplane spotter to claim knowledge that a plane is an F even though he cannot eliminate the unlikely possibility that the plane is a G. But, as I argued earlier, this knowledge claim will *not* be appropriate if the spotter knows about the existence of G's. This fact strongly suggests that practical conditions do not lower the standards for appropriate knowledge attribution. For without any shift in the practical circumstances (either in the speaker's goals or in the practical situation), a possibility — such as the possibility that a plane is a G — can become relevant to the appropriateness of a particular knowledge claim in virtue of a shift in the evidence which the speaker recognizes.

Likewise, practical considerations do not *widen* the range of relevant error possibilities. Imagine a situation in which there is no interest whatsoever in anything other than the truth of the matter under consideration and no significant practical constraint on one's ability to gather the relevant information. Is it appropriate, in such circumstances, to claim to know (for example) that a certain bird is a goldfinch, even if one has not dissected it in order to insure that it is not an alien spying contraption, checked with local genetic engineers to establish that it isn't a modified bluebird, or searched the surroundings for signs of a holographic apparatus? Of course it is appropriate. Regardless of one's purposes and practical circumstances, one does not have to attempt to acquire detailed and specific evidence against these possibilities. To do so would be neurotic at best. Admittedly, if one had evidence suggesting that trickery, intergalactic subterfuge, or genetic manipulation might be involved, then such possibilities would be relevant. But that is not a practical matter. It is a matter of what is supported by the reasons in one's possession.

It is true that when the stakes are high there is some tendency to withhold knowledge claims even if one possesses evidence which one would ordinarily regard as warranting claiming knowledge. However, this behavior can be satisfactorily explained

without Practical Constraints Contextualism. For one thing, not all of our attributive behavior is reasonable in the relevant sense. Stewart Cohen offers an example in which a person, Mary, refuses to accept that someone else, Smith, knows on the basis of his printed itinerary that a certain flight will stop in Chicago, even though neither party has any particular reason to think that his itinerary is incorrect. Cohen urges that this response is quite reasonable, since Mary is concerned to meet an important business contact in Chicago.¹⁹ It seems to me, however, that Mary's behavior is unreasonable, though perfectly understandable. When we are fearful and anxious we tend to lose confidence in the truth of our beliefs and consequently are hesitant to claim knowledge; we check and recheck, worrying about extremely unlikely possibilities of error, even though we have no reason to suspect that the claim in question is false. Thus in conditions of practical extremity we might lose sight of the force of the reasons of which we are aware and judge, unreasonably, that the person in question lacks knowledge. Such considerations also provide an adequate explanation of people's (supposed) hesitation to claim knowledge when queried in the courtroom. (There may also be an epistemic explanation, at least in some cases, since the lawyer's question, "Do you really know that ___?", can suggest that he knows something you don't know. The mere fact that he has asked the question may seem to provide a reason to doubt the truth of your belief.)

Our attributive practice may also be affected by semantic context-sensitivity in the embedded sentences stating *what* we know (or don't know), and this may make it seem that high stakes raise the standards. For instance, we sometimes refrain from claiming knowledge because we discover that the practical circumstances demand a degree of precision and exactness higher than that normally called for. If you believe your watch to be working properly, you would ordinarily claim knowledge that it is five o'clock on the basis of what your watch says. But you might not do so if you know that your hearer is

¹⁹ "Contextualism and Skepticism," pp. 95 ff.

engaged in a sensitive experiment and needs to know whether it is four fifty-nine and fifty-nine seconds or five o'clock on the dot. Examples of this sort do not indicate that higher standards for appropriate knowledge attribution are in place, because what shifts is not the standard for appropriate knowledge attribution, but rather the issue under consideration (though it is still expressed by the same form of words).

It seems, then, that Practical Constraints Contextualism does not have the right machinery for explaining shifts in the standards for appropriate knowledge attribution. What matters is not the practical situation but rather one's understanding of the *epistemic* situation. It is appropriate to attribute knowledge to someone just when you take that person to have decisive, specific evidence against those possibilities of error *which you take to have some reason in their favor*. If you take there to be no reason to suspect that a certain error possibility might obtain, then it would be unreasonable, regardless of the practical setting, to deny that someone has knowledge on account of the fact that he or she does not have specific evidence against that possibility. Likewise, if you take there to be good reason to suspect that a certain error possibility might obtain, then it would be inappropriate, regardless of the practical setting, to attribute knowledge to a given person unless you thought that he or she had decisive, specific evidence against the possibility in question. Shifts in our understanding of the evidential situation are thus the primary source of shifts in the set of possibilities which we deem it appropriate to bring to bear in assessing a person's knowledge. What is pivotal here are not practical considerations but epistemic reasons — reasons for doubt and for belief.

What it is for a possibility to have some reason in its favor or for a person to have "decisive, specific evidence" against a possibility's obtaining? A developed theory of these matters would be desirable, but it is not necessary for my purposes here. We have a perfectly good understanding of these notions, manifested in our practices. For instance, if you are looking at what you take to be a goldfinch in ordinary circumstances, you will take there to be some (prima facie) reason to suspect that it might be a similar-looking bird of

another sort (though this reason may be defeated by your evidence). But you will take there to be no reason at all as things currently stand to suspect that it is an alien spying device. Likewise, suppose that you know that local genetic engineers have recently released modified bluebirds that look like goldfinches. Then, in contrast with the ordinary case, you will think it inappropriate to say that a local expert birdwatcher knows that he is looking at a goldfinch unless you believe (1) that he knows about a mark which distinguishes goldfinches from the modified bluebirds and (2) that he has recognized this mark in the particular case. Recognizing such a mark would count as having "decisive, specific evidence" against the possibility in question. My aim here is only to point out the significance and role of these notions in our practices of knowledge attribution; this role is a datum which any adequate theory would have to respect.

b. The Topic of Epistemic Evaluation

What are we assessing about a person when we decide whether to attribute knowledge to him or her? According to Practical Constraints Contextualism, the goal of epistemic assessment is to determine whether the range of error possibilities that the person is in a position to rule out is close enough to the ideal position of being able to rule out all possibilities of error. On this view, the strength of one's epistemic position is determined by the range of error possibilities which one can rule out. The wider the range, the better the position.

This view is not borne out by our actual evaluative practice. Suppose that I arrive at home, place the grocery bags on the kitchen table, and turn to leave the room. If asked, I would quite appropriately claim that I know where the groceries are. However, because I cannot currently see them, there are any number of possibilities which are compatible with my current evidence. For instance, it is possible, so far as my evidence goes, that as soon as my back was turned a cleverly designed pneumatic trap-door silently eased the whole table, with the groceries upon it, into the basement. I could obtain strong evidence against

this possibility simply by turning around and checking. But would doing so constitute any improvement in my epistemic position, given my actual circumstances? If circumstances were different and there were some reason for me to think that such a thing had occurred, then my epistemic position would be improved by turning around to check. But if I judge that there is no reason even to suspect that anything untoward has occurred, then I also judge that my epistemic position with regard to the location of the groceries is already as good as it can be. I judge that there is nothing that I need to do to improve it, no evidence that I need to overcome or explain away. I am not merely judging that it is good enough for my practical purposes. I am judging that it couldn't be better, given my actual circumstances. That is why I am prepared to claim knowledge.

It is true that I might nonetheless lose my confidence that the groceries are on the table. If I did, a quick glance would be an excellent way to assuage the worry. But this does not indicate that my epistemic position would be improved if I were now to check. Repeated checking is not always epistemic progress. Consider the case of the worrywart who carefully locks the front door, has a memory of doing so, and has no reason to think that he didn't. As he drives away, he begins to worry. Going back and checking might make him feel better, but it would not improve his epistemic position. He already has every reason to believe that he locked the door. His position with respect to that issue is as good as it can be; he just doesn't accept that it is. So the mere fact that one can feel worried does not show that one's epistemic position stands in need of improvement. What matters is rather whether there is reason for worry.

In sum, then, our attributive practices treat the *strength* or *goodness* of someone's epistemic position as being determined by the extent to which the person possesses decisive, specific evidence against those possibilities of error which there is some reason to believe or suspect to be the case. If one has decisive evidence against all such possibilities, then one is in the strongest or best epistemic position; nothing is gained by acquiring additional evidence against alternatives which one already recognizes have no reasons in

their favor. Thus our epistemic evaluations are not based on a determination of whether the person's epistemic position is close enough, for practical purposes, to the ideal of being able to rule out every possible alternative. Practical Constraints Contextualism is a faulty account of our practices of epistemic assessment because it flies in the face of the fact that judgments about *what there is reason to believe* guide our epistemic evaluations, not anything having to do with practical considerations at all.

We are now in possession of a positive argument against Infallibilist skepticism. Without adopting the Practical Constraints View, Infallibilist skepticism has no plausible explanation of our ordinary practices of knowledge attribution. But the Practical Constraints View ultimately yields an incorrect account of the appropriateness conditions of knowledge attributions. Consequently, we have good reason to reject the Infallibilist's explanation of our attributive practice and to continue to take our ordinary knowledge assessments at face value. Since we often appropriately judge that people possess knowledge even though they cannot meet the Infallibility Requirement, we should conclude that this requirement is not correct. Knowledge does not require infallibility.

One way to put this result is to say that the Infallibilist has misunderstood or misconceived the requirements that one must meet in order for it to be true to say that one possesses knowledge. However, this can suggest that we have gained only a verbal advantage over the Infallibilist skeptic. For one can imagine the Infallibilist saying, "You've merely shown that my point is poorly stated if it is put as a claim about knowledge. My fundamental point is that since the best or ideal epistemic position is being able to rule out all possible ways in which we could be wrong, we can never attain the best or ideal epistemic position. What we call 'knowledge' is merely second-best. Our fallibility is a shortcoming and a cause for disappointment."

We would have attained nothing but a verbal victory if meeting the Infallibility Requirement figured as the ideal in our practices of knowledge assessment. For then infallibility would be our touchstone, the position against which we measure ourselves and

inevitably fall short. Thus while we might be content with our epistemic capacities for the purposes of everyday life, we would be doomed to disappointment when we reflected on them from a purely epistemological standpoint. This would remain the case even if we disagreed with the Infallibilist over whether *knowledge* requires infallibility. This problem plagues contextualist responses to Infallibilist skepticism, as Lewis admits: "Never —well, hardly ever — does our knowledge rest entirely on elimination and not at all on ignoring. *So hardly ever is it quite as good as we might wish.* To that extent, the lesson of skepticism is right — and right permanently, not just in the temporary and special context of epistemology".²⁰ Fortunately, however, infallibility does not figure even as the ideal in our ordinary practices of knowledge assessment. When we assess people's knowledge, our concern is to determine whether they have decisive evidence against those specific possibilities of error which have some reason in their favor. We want to know how they stand in relation to the reasons that there are for and against the truth of their beliefs. Even the best possible position of this sort does not yield infallibly true beliefs. Consequently, infallibility is not the touchstone of our ordinary knowledge assessments, and our fallibility is not a reason for epistemic disappointment.

4. The Prospects for Skepticism

Our investigation of Infallibilist skepticism yields a more general lesson. Since our practices of knowledge assessment are not responsive to merely practical or conversational concerns, we have every right to take them to reveal the requirements for knowledge possession. Hence, we can reasonably reject any form of skepticism which, like Infallibilist Skepticism, imposes requirements to which we are not committed by our ordinary epistemic practices.

²⁰ "Elusive Knowledge," p.563, italics added.

This result does not completely vanquish skepticism. Since the appropriateness of our ordinary knowledge attributions does not entail their truth, it is possible that we are wrong, though warranted, in thinking that the conditions for possessing knowledge are ever met. However, our investigation has clarified how skepticism could turn out to be true: it won't be true unless our ordinary knowledge attributions amount to epistemically reasonable errors. In particular, we must be making a reasonable error in thinking that we ever meet the requirements which we insist upon in ordinary life, or we must be making a reasonable error in thinking that only the requirements which we ordinarily recognize must be met. As I noted earlier, however, it will do no good for the skeptic to insist that the concept of knowledge, quite apart from our epistemic practices, involves requirements of which most people — masters of our language and epistemic practices — are simply ignorant. For given the failure of the Practical Constraints View, the skeptic would then be left saying, implausibly, that our ordinary practices are just *mistaken*. Through what kind of special insight has the skeptic discovered this fact, and why have the rest of us missed the boat? That approach will never fly. The skeptic consequently needs to find his unmeetable conditions for knowledge possession within the domain of our ordinary – and fallibilist — epistemic practices. He needs to establish that those practices reveal our commitment, or provide us with reason to be committed, to the epistemic requirements which he insists that we must meet. And he needs to show how we could reasonably have been unaware that we must meet these requirements, or — if we are aware that we must meet them — how we could reasonably think, incorrectly, that we do.

Can these things be shown? In order to answer this question, we need a deeper understanding of the requirements involved in our ordinary practices of knowledge

evaluation. And that, in turn, requires a careful and detailed investigation of our actual epistemic lives.²¹

Word count: 9749 (including this note)

²¹ I would like to thank Paul Eisenberg, Warren Goldfarb, Steven Gross, Tom Kelly, Dick Moran, Jim Pryor, Angie Smith and Barry Stroud for critical comments and helpful advice.