

Epistemic Gradualism and Ordinary Epistemic Practice: Response to Hetherington

Adam Leite

Published online: 22 November 2006
© Springer Science + Business Media B.V. 2006

Abstract This paper responds to Stephen Hetherington's discussion of my 'Is Fallibility an Epistemological Shortcoming?' (2004). The Infallibilist skeptic holds that in order to know something, one must be able to rule out every possible alternative to the truth of one's belief. This requirement is false. In this paper I first clarify this requirement's relation to our ordinary practice. I then turn to a more fundamental issue. The Infallibilist holds – along with many non-skeptical epistemologists – that Infallibility is *epistemically* superior to the epistemic position attained when we have (what we ordinarily call) knowledge. This is false, too, as our ordinary practices show. Ordinary epistemic appraisal does not concern our standing on a scale of evaluation which has Infallibility at its apex. For this reason, even if gradualism is correct, it does not show how Infallibilist skepticism can arise out of our ordinary practice.

Keywords infallibility · epistemic gradualism · skepticism · knowledge · fallibilism

Skepticism and Ordinary Practice

“No one knows anything about the world around us,” the skeptic about knowledge of the external world declares. Taken at face value, this claim appears to be false. We therefore need to ask two questions about it. First, does it conflict with what we say in the context of an ordinary conversation when we say such things as, “I know where my wife's keys are”? Second, is there any reason to think that it is true? If the answer to both questions is, “Yes,” then we have a serious skeptical threat on our hands — an argument which threatens to undermine much of what we say and do in ordinary life. If the skeptic's claim conflicts with our ordinary knowledge claims but there is no reason to think it true, then it can be dismissed summarily; it will be of no more interest than the claim, for example, that George W. Bush is a woman. If the two claims do not conflict, however, then the situation is interestingly more complicated. Everything will depend upon what the skeptic's claim means. It may be that the skeptic is merely engaged in outrageously changing the subject, like Paul Edwards' fanciful character who declares, “There are no doctors in New York,” but turns out to mean that there is no one in New York who can cure any possible illness in under two minutes.¹ In this case, the skeptic may safely be ignored.

¹Paul Edwards, “Bertrand Russell's Doubts about Induction,” *Logic and Language*, first series, Antony Flew, ed., Oxford: Basil Blackwell, 1951.

A. Leite (✉)
Department of Philosophy, Indiana University, Bloomington, IN 47405, USA
e-mail: aleite@indiana.edu

Or it may be that the skeptic is making an interesting and important point, but not saying exactly what – at first blush – he seems to be saying. In that case, we would do well to give skepticism a serious hearing, not because it threatens what we do and say in ordinary life (it doesn't), but because it has something to teach us.

My concern in “Is Fallibility an Epistemological Shortcoming?” was to assess a particular skeptical argument in the above terms.² The argument depends upon the *Infallibility Requirement*, the requirement that in order to know any *p* one must be able to rule out or eliminate every possible alternative to the truth of *p* (where “rule out” and “eliminate” are interpreted in such a way that satisfying this requirement would put one in a position in which it is not possible for one's belief, held on the basis on which it is held, to be false). My question was whether this argument provides a good reason to believe a skeptical conclusion which conflicts with our ordinary knowledge claims. I maintained that it does not. Here's why. It is *prima facie* reasonable to reject the Infallibility Requirement, since that requirement plays no role in our ordinary practices of epistemic assessment. So if one insists upon the Infallibility Requirement anyway, one must provide a plausible explanation of why it doesn't figure in our ordinary epistemic practices. According to the best attempt in this direction, we ordinarily ignore the Infallibility requirement (even though knowledge requires satisfaction of it) *because of the demands of the practical and conversational situations in which we ordinarily decide whether people know things*. For this reason, it is suggested, considerations about when it is ordinarily appropriate to attribute knowledge don't by themselves reveal the conditions one must meet in order to possess knowledge. This “Practical Constraints View,” as I call it, fails because it yields incorrect predictions about our ordinary practices of knowledge attribution. In particular, as I argue, it yields incorrect predictions about our judgments regarding when it is ordinarily appropriate to attribute or deny knowledge. We are therefore free to take our ordinary practices of epistemic assessment at face value: knowledge does not require satisfaction of the Infallibility Requirement.

Stephen Hetherington³ doesn't dispute this argument, and I will presuppose it in the following discussion. Where Hetherington disagrees with me is over its significance. On his view, the Infallibilist skeptic tells us something correct about our epistemic situation, even if what this skeptic says doesn't quite conflict with what we say in ordinary life when we claim to know things. And on his view what the Infallibilist skeptic tells us is something that is *importantly related* to these ordinary knowledge claims. As we might put it, what this form of skepticism reveals is that we don't know anything as well as we ideally might. I disagree with Hetherington on these points. My aim here is to elaborate some of what I said in my paper in order to explain why. I will first clarify the relation between my proposals and the view which Hetherington calls “challenge conservatism.” I will then clarify my conception of the relation between the Infallibility Requirement and the requirements which govern our ordinary practices of knowledge attribution. Finally, I will argue against a crucial assumption which Hetherington and the Infallibilist skeptic share: the assumption that the position of Infallibility is *epistemically superior* to the position which we are content to call “knowledge” in ordinary life. This argument forms the heart of this paper. The assumption underlies a great deal of epistemological reflection. But it is false.

² “Is Fallibility an Epistemological Shortcoming?” *The Philosophical Quarterly*, 54: 215, April 2004, pp. 232–251.

³ “Scepticism and Ordinary Epistemic Practice”, *Philosophia*, 34: 3, 2006, pp. 303–310.

Ordinary Practices of Knowledge Assessment and the Conduct of Inquiry

According to my account of our ordinary practices of knowledge attribution,

It is appropriate to attribute knowledge to someone when you take that person to have decisive specific evidence against those possibilities of error *which you take to have some reason in their favour*. If you take there to be no reason to suspect that there is any possibility of a certain error, then it would be unreasonable, regardless of the practical setting, to deny knowledge to people on account of the fact that they do not have specific evidence against that possibility. Likewise, if you take there to be good reason to suspect that there is a possibility of a certain error, then it would be inappropriate, regardless of the practical setting, to attribute knowledge to people unless you thought that they had decisive specific evidence against the possibility in question (p. 247).⁴

It should be stressed that these claims are meant to characterize features of our *attributive practice*, in particular the feature(s) of the person that we consider when we determine whether or not a person knows some given thing and the conditions under which we regard it as appropriate or inappropriate to attribute knowledge to a person. These claims are not themselves claims about the conditions one must meet in order to possess knowledge, nor do they directly entail any such claims. It is true that I ultimately draw some conclusions about the requirements for knowledge from these considerations. In particular, I ultimately take them to provide good reason for rejecting the Infallibility Requirement. But beyond this I do not explicitly draw any further, positive conclusions about the requirements for knowledge, nor need I do so for the purposes of refuting Infallibilist skepticism.

In particular, the above quoted claims do not commit me to the view which Hetherington calls “challenge conservatism”, according to which “Ordinary epistemic practice requires us (if we are to have some particular piece of knowledge) to eliminate only possibilities that our existing reasons or evidence *already support to some extent*” (p. 304). Moreover, as a characterization of the conditions for *knowledge*, challenge conservatism is arguably false. Consider, for instance, Goldman’s example of Henry who sees a barn surrounded by barn facades which are visually indistinguishable from real barns from the road.⁵ Henry does not know about the existence of the barn facades, so he does not possess any reasons or evidence which support the possibility that what he sees is only a barn façade. Still, he does not know that what he sees is a real barn; to know that, he would (at least) need to be able to eliminate the possibility that it is just a barn façade. To put the point more generally and in the idiom I favor, in order to know something, one must be able to muster decisive positive evidence against those possibilities of error which *have some reason in their favor*. There can, of

⁴ The published version of this passage, reprinted above, does not faithfully reflect my original manuscript because of errors introduced during the copy-editing process. It should read (with corrections in italics),

... If you take there to be no reason to suspect *that a certain possibility of error obtains*, then it would be unreasonable, regardless of the practical setting, to deny knowledge to people on account of the fact that they do not have specific evidence against that possibility. Likewise, if you take there to be good reason to suspect *that a certain possibility of error obtains*, then it would be inappropriate, regardless of the practical setting, to attribute knowledge to people unless you thought that they had decisive specific evidence against the possibility in question.

By “possibility of error,” I mean any possibility which is incompatible with *p* or with the person’s knowing that *p*.

⁵ Goldman, Alvin. ‘Discrimination and Perceptual Knowledge,’ in *Essays on Knowledge and Justification*, Pappas and Swain, eds., Cornell University Press, 1978, pp. 120–145. The example is reported to be from Carl Ginet.

course, be reasons of which one is unaware; what possibilities of error have some reason in their favor depends upon *what else is the case* in the *objective* circumstances.

Still, I do accept a view analogous to “challenge conservatism” as an account of our practices of epistemic evaluation and attribution (see the quote above). Hetherington charges that this sort of account is at odds with our ordinary practices of inquiry, for three reasons. First, we sometimes realize that we do not possess the best evidence with which to investigate some issue, and so – if forced to investigate it nonetheless – will reflect upon possibilities for which our evidence provides no support (p. 304). Second, this willingness appears to be a feature of ordinary inquiry more generally: “ordinary practice might often incorporate a greater chance of people being open to new possibilities being thought of and introduced for testing – without there already having been any independent and prior establishing of the bona fides of those new ideas” (p. 304). Finally, Hetherington charges that adhering to challenge conservatism would lead us to miss out on the potential epistemic benefits of open-mindedness to possibilities which are not supported by our evidence, since such open-mindedness can result in “new lines of inquiry luckily setting the ‘ordinary’ inquirers on the path to gaining new knowledge which they would not have discovered if they had allowed themselves to ignore whatever was dismissed as being irrelevant by their existing evidence” (pp. 304–305).

To begin with the first two points, I do not think that ordinary inquiry is as open to investigation of unsupported alternatives as Hetherington suggests. Suppose that you and I are peering through a window overlooking my back yard in order to determine whether a goldfinch is at the bird feeder. Some possibilities which you might suggest will be regarded as serious and helpful contributions to the inquiry. But if you suggest that what’s at the feeder might be an intergalactic spying device, I and any other sane, reasonably well-informed auditor will regard your suggestion as useless or facetious. I suggest that this is because we take there to be no reason whatsoever to suspect that it might be an intergalactic spying device. It seems to me that this example generalizes. Even in everyday, nontechnical inquiry, we limit our serious attention to those possibilities which we take to have some reason in their favor. Admittedly, in certain situations we recognize that our current evidential position isn’t especially good and so cast our nets fairly widely. This is often the case at the beginning of an inquiry when we are largely ignorant of the subject matter. But even here, our choice of lines of inquiry is guided by background assumptions and principles and our understanding of what they favor or support. (Of course, it is also part of ordinary inquiry that these background considerations can reasonably be modified in the course of inquiry as new evidence emerges.) At the outer limit there might be situations in which we are only in a position to say that so far as we can see, nothing prevents any of a certain set of alternatives from being the case. Here it becomes apt to say that we have reason to suspect, for each of these alternatives, that it is the case — more reason, at least, than we have for those alternatives which are already off the table because there is decisive reason to believe they do not obtain. Likewise, we sometimes judge that our background evidence is not sufficient for us to make a reasonable judgment about whether or not a certain possibility has some reason in its favor. In such a case, we reasonably seek to acquire further information. But here, too, we are guided by our background understanding of what is the case and what those things support or tell against. And crucially, if we judge that there really is no reason whatsoever in favor of a certain possibility, we give it no further thought — unless we find reason to doubt that judgment.

Let us now consider Hetherington’s last point. If our conduct of inquiry were guided by our sense of what the reasons support, would we lose potential benefits? Consider a strategy of inquiry which systematically sought specific decisive evidence for or against *every logically possible alternative*. Given the finitude of our life spans and resources, the costs of

such a strategy would outweigh its benefits. We would never answer a single question. So it seems to me that a reasonable approach to inquiry is to guide our selection of alternatives for investigation by our best current sense of what the reasons support. Nonetheless, there is an important question here. It is an empirical question. Do we attain better results in inquiry if we are guided by our sense of what the reasons support? Might we do better if we simply investigated whatever possibilities came into our minds or followed some other non-evidence-based strategy? It seems to me that we wouldn't. But beyond anecdotal considerations, I don't have empirical evidence to back that judgment up.

However this might be, there is a more fundamental issue here. My account of our practices of knowledge assessment and attribution doesn't in fact directly dictate anything about acceptable patterns of inquiry. To see this, consider an example. I believe that it is now known that HIV is the cause of AIDS, and that there is no good reason to suspect that anything else is the cause. Accordingly, if I were an AIDS researcher, I would spend my time seeking cures and not exploring other possible causes of AIDS. It is perfectly compatible with these commitments that I think it proper that some (few) AIDS researchers continue to be funded despite the fact that they are investigating other possible causes such as drug use and environmental factors. After all, I recognize that I am a fallible being. I may have misread the data. Given the resources available and the importance of the subject, I therefore endorse a social, institutional practice of inquiry in which some (few) researchers continue to pursue what are widely accepted as dead ends. Of course, as things stand, we expect such researchers to be able to provide some rationalization of their research programs, some account of where the common thinking goes wrong and why they expect their research to bear fruit. But even this is not demanded by my account of our practices of knowledge assessment and attribution. If we had more resources and there were empirical evidence favoring the potential fruitfulness of such an arrangement, we might even accept – compatibly with my account – institutional arrangements in which a certain percentage of researchers were permitted to investigate whatever possibilities came into their minds, or were indicated by that week's horoscope, or whatever. These sorts of institutions of inquiry would be compatible with my account because like all else we do, even our knowledge attributions are fallible. We might judge that in some cases such researchers are investigating possibilities which are known not to obtain. But recognizing the fallibility of such judgments, we might deem it worthwhile to deploy whatever resources seem helpful in providing a further check. Considerations about how we proceed in inquiry thus do not tell against my account of our ordinary practices of knowledge attribution.

The Skeptic

Hetherington casts the skeptic as challenging us to be more open-minded in order to obtain an even *better* epistemic position than that deemed knowledge by our ordinary practices of epistemic assessment. The skeptic, he suggests, is inviting us to seek “extraordinary knowledge”, not mere “ordinary knowledge,” where “extraordinary knowledge” is something which lies on the same scale of epistemic assessment as “ordinary knowledge”, but is *better*. The skeptic does this, Hetherington suggests, by inviting us to consider possibilities, such as the possibility that we are dreaming, or that we are just brains in vats, which we had not previously considered and were not required to know not to obtain. Here, Hetherington is working with the following conception. Ordinary knowledge that p is a matter of being able to eliminate some select subgroup of the possible alternatives to the truth of p . The more possible alternatives one can eliminate, the more extraordinary one's

knowledge becomes. At the limit, one would be able to eliminate all of them. That's the best, the most extraordinary, epistemic position. Hetherington thus accepts a version of a Relevant Alternatives account of ordinary knowledge. And like the paradigm Relevant Alternatives theorists (such as Dretske and Nozick), Hetherington denies that (ordinary) knowledge is closed under known implication.⁶

One could accept my account of our practices of knowledge attribution without conceiving of "ordinary knowledge" in this way. For instance, here is a natural view.

In order to know (in the ordinary sense of that word) that I am sitting at my desk as I write these words, I need to know that I am not a brain in a vat and that I am not dreaming right now. Moreover, I *do* know these things. But this knowledge is fallible, inasmuch as its basis does not logically or metaphysically guarantee that I am not a brain in a vat or just dreaming. So I don't meet the Infallibility Requirement. But I *do* satisfy the closure principle.

There doesn't seem to be any reason in principle why this conception of the situation couldn't be right. However, one could understand the situation in this way while accepting my account of our practice. For one could hold that we don't need to muster specific evidence against these possibilities because we recognize that there is no reason to suspect that they obtain.

I myself favor a view of roughly this sort. For this reason, I think that Hetherington's comparison of "ordinary" and "extraordinary" knowledge misses something important, insofar as it turns upon the idea that the skeptic is inviting us to bring in more and more possible alternatives. The shift to the Infallibilist's "extraordinary" demands isn't brought about mainly by bringing in alternatives which we previously were not required to know not to obtain; rather, it involves a shift in what is to count as "eliminating" or "ruling out" an alternative, where the shift is away from all that we can have when it comes to contingent claims about the world: fallible reasons for thinking that a possibility does or does not obtain.

This point returns us to the question of the relationship between the Infallibilist skeptic's utterance, "No one knows anything about the world," and the knowledge claims we make in ordinary life. I don't deny that someone might reasonably decide to engage in an extraordinary form of inquiry, governed by unusual and extremely stringent methodological rules (e.g., a demand for infallibility). I could even be convinced that something intellectually useful might come of such an enterprise. But if those methodological rules do not govern or appear in our ordinary epistemic practice, then even if the upshot of the investigation is negative, it would be misleading to state that upshot by saying, e.g., "We don't know anything in the relevant domain." And it would likewise be invalid to conclude from such an upshot that we don't know the things in question.

Scales of Epistemic Assessment: What is our Ordinary Epistemic Ideal?

Hetherington's "gradualism" threatens to dampen the force of the preceding considerations. Gradualism is the view that knowledge comes in degrees; some instances of knowing that *p* are better, qua knowledge, than others, with *Infallibility* being the top of the scale. If this is right, then my argument against the Infallibility Requirement seems to show only that satisfaction

⁶ *Good Knowledge, Bad Knowledge*, Oxford University Press, New York, NY: 2002, chapter 2.

of the Infallibility Requirement isn't required for the kind of knowledge that we are content with in ordinary life. And if the gradualist is right that that kind of knowledge lies on a scale with Infallibility at its top, then while the skeptic shouldn't conclude, bluntly, that we don't know anything about the world, it would be correct for the skeptic to conclude that we don't know anything about the world *as well as we'd like or as well as would be ideal*.

We can distinguish two aspects of Hetherington's gradualism. First, the schematic idea that knowledge comes in degrees; second, the substantive claim that these degrees lie on a scale which has Infallibility at its top. I have qualms about the first idea, since it seems to me that the linguistic data do not support it. (For instance, the relevant comparative formulations do not roll naturally off the tongue, and without theory-driven elaboration it is quite unclear what they might mean. In this regard, I side with Dretske's quip, "[Knowledge] is like being pregnant: an all or nothing affair".⁷) However, I do not want to press these issues now. For the second claim of Hetherington's gradualism reflects a deeper thought which informs much philosophical reflection in these areas. This is the thought that attaining the position of Infallibility would be something recognizably like knowing – something lying on the very same scale of appraisal – only better. Or, to put it another way, that Infallibility is just like knowledge, *only more so*. As Hetherington puts it, "extraordinary knowledge that *p* is epistemically better, all else being equal, than is ordinary knowledge that *p*" (p. [5]). This underlying thought does not require Hetherington's view that knowledge comes in degrees, or even the view that there are (at least) two kinds of knowledge, ordinary and extraordinary. All it requires is that there be a scale of epistemic appraisal on which epistemic positions can be ranked, and that on this scale the position we ordinarily call 'knowledge' occupies a lower rank than the position of Infallibility. As long as this thought remains unaddressed, any refutation of Infallibilism will seem dissatisfying. It will seem that by showing that *knowledge* doesn't require Infallibility, we just show that in ordinary life we are content to make do — to settle for second best or worse. It will accordingly seem, as it seems to Hetherington, that from a purely epistemological standpoint, we aren't as well off in ordinary life as we might ideally be. As I will argue, however, the thought underlying this dissatisfaction is incorrect, as our ordinary practices of knowledge attribution show.

Before we begin, it should be noted that my strategy here is compatible with accepting the first claim of Hetherington's gradualism, viz., the schematic idea that knowledge comes in degrees. For this reason, Hetherington is wrong to charge me with sharing the skeptic's commitment to "qualitative absolutism", the view that "all possible instances of knowledge that *p* share only one possible epistemic quality ... *as cases of knowledge that p*" (p. [308]). For my purposes here (and in my original paper) I don't *have* to share this commitment, though in

⁷ "The Pragmatic Dimension of Knowledge," p. 48, Dretske, *Perception, Knowledge, and Belief*, New York: Cambridge University Press, 2000. Certain common expressions might seem to favor the gradualist approach. For instance, we commonly say such things as, "I don't know her as well as he does," "Student X knows modal logic better now that he's finally taken a course in it," and "John knows how to play chess well." However, these examples – and all plausible examples of which I am aware – concern what we might call *acquaintance, know-how, understanding* and other forms of knowledge whose relation to propositional knowledge is quite unclear. I am not aware of any example of ordinary usage which clearly supports the gradualist account of propositional knowledge. (Hetherington's examples in this regard strike me as forced and unnatural. Consider, for instance, "I do know that I locked my office door. It is true, though, that if I go back to check on whether I locked it, I will improve that knowledge slightly" (*Good Knowledge, Bad Knowledge*, p. 2).) For these reasons, the linguistic data are not probative, unless it can be shown that acquaintance, know-how and the rest all depend upon propositional knowledge in such a way that the gradualism which they display must be regarded as *derivative from* gradualism in the underlying propositional knowledge. In addition, the case will not be closed until some account is given of why we do not commonly talk about propositional knowledge in the way gradualism would predict.

fact I incline towards it; what I reject is the idea that if knowledge comes in degrees, those degrees lie on a scale of evaluation which has the position of Infallibility at its apex.

To keep things straight in what follows, I will use the term ‘knowledge’ for what Hetherington means to pick out by the phrase ‘ordinary knowledge’. I will use ‘Infallibility’ to denote the position of being able to rule out every possible alternative (where, again, “rule out” is interpreted in such a way that this position entails infallibly true belief). It should be kept in mind, too, that for our purposes here skepticism – understood as the view that no one knows anything about the world because we cannot meet the Infallibility requirement – is already off the table. We are taking for granted that we can attain knowledge in the course of our ordinary lives, and that our ordinary practices of epistemic assessment and knowledge attribution are a good guide to the conditions one must meet in order to know something.⁸ Our question – to be answered in the light of these assumptions – is just this: Is Infallibility epistemically superior to knowledge?

An analogy may help one see how a negative answer could be possible here. For the purposes of playing basketball, it is better to be seven feet tall than to be six feet tall. In this case, an increase in height constitutes an improvement. But more is not always better. Given the world as it is, it would not be better for the purposes of playing basketball to be twenty-five feet tall. In fact, it would be detrimental. Consequently, our appraisals of basketball players’ heights are not governed by a scale of appraisal which has *being as tall as is logically and metaphysically possible* at its apex.⁹ For this analogy to serve its purpose, it need not be claimed that Infallibility would be epistemically detrimental. Rather, what the analogy reveals is that even though a certain shift along a given dimension may be an improvement relative to a particular scale of evaluation, further shifts along that dimension may not be improvement at all relative to that scale of evaluation. My suggestion, then, is that something similar goes on in the epistemic realm: Infallibility is not any improvement, relative to the scale of evaluation that is relevant to epistemic assessment. This isn’t to say that Infallibility would be epistemically inferior. It would just be different.

Before I can defend this suggestion, I need to clarify some important terms. First, “epistemic position”.

Necessarily, if one knows that p , then p is true. Consequently, if we think of *knowledge* as an epistemic position, then we have to agree that it is an infallible epistemic position: one cannot be in it with regard to any proposition unless that proposition is true. Interpreting the term ‘epistemic position’ in such a way as to include knowledge therefore trivially precludes a view like Hetherington’s: on the assumption that infallibly true belief is the ideal, Infallibility (as defined above) would not be a superior epistemic position to knowledge, because knowledge, too, would be an infallible epistemic position. To give substance to our overall question, then, we have to think of knowledge as a state involving two factors: (1) the person’s epistemic position with regard to p , and (2) the fact as to whether or not p , where it is understood that for at least some epistemic positions adequate for knowledge of a given p , one can be in those epistemic positions regardless of whether or not p is true.¹⁰ In these terms,

⁸ Establishing these points is the main burden of “Is Fallibility an Epistemological Shortcoming?”. Hetherington does not disagree with this aspect of my argument.

⁹ I am grateful to Ram Neta for this example.

¹⁰ This conception of knowledge is disputed by Williamson (*Knowledge and Its Limits*, Oxford University Press, 2000, esp. chapters 1–3) and McDowell (“Knowledge and the Internal,” in *Meaning, Knowledge and Reality*, Cambridge, MA: Harvard University Press, 1998). In my “On Williamson’s Arguments that Knowledge is a Mental State” (*Ratio (new series)*, vol. XVIII no. 2, June 2005, pp. 165–75), I argue that Williamson’s arguments against this conception fail.

Infallibility (as defined above) is a position which logically or metaphysically entails the truth of the believed proposition, so that necessarily, if one is in that position regarding p , then p is true. By contrast, the epistemic position involved in knowledge is fallible, in this sense: when one knows that p , one's epistemic position with regard to p is logically and metaphysically compatible with p 's falsehood. Or, to put it another way, it is possible for one to be in just the epistemic position which one is in when one knows p , and yet for p to be false. Putting our overall question in these terms, the question is whether Infallibility is epistemically superior to the epistemic position which is involved in knowledge.

Epistemic positions can be understood as involving a number of elements, such as the person's evidence, how that evidence was acquired, how the belief was formed, or the person's causal or modal relation to the facts as to whether p . What is crucial for our discussion here, however, is a conception of epistemic positions as being determined primarily by the evidence possessed by the person and its relation to the alternatives to the truth of one's belief. This is how Hetherington conceives of the matter,¹¹ and the notion of "ruling out" is generally understood in evidentialist terms as well. Since I believe in any case that an evidentialist account of knowledge is broadly right, I will follow Hetherington in this regard.

A second issue crucially requiring clarification is what is meant by the phrase 'epistemically superior'. For there are a variety of ways in which Infallibility might be regarded as better than the epistemic position involved in knowledge. For instance, Infallibility would be *instrumentally* desirable, insofar as it would reduce the chance of failure in our attempts to achieve our goals. But that isn't what is at issue here. And it would do no good simply to ask without qualification whether Infallibility is better or superior to the epistemic position involved in knowledge, since it is quite unclear how this question should be answered. (Would it matter, for instance, if one simply preferred the thrill of running a slight risk?) The term 'epistemic' is thus essential here, and the question is accordingly whether Infallibility occupies a superior position on the *epistemic* scale of evaluation. To answer this question, we must determine which scale of evaluation that is and what position the epistemic position involved in knowledge occupies on it.

'Epistemic' is an epistemologist's term of art, open to interpretation and even definitional fiat. But any interpretation or definition is subject to an important constraint: the epistemic has something to do with knowledge; the position which we occupy when we have knowledge is an epistemic position if anything is, and the kind of assessment we engage in when we decide in ordinary life whether people know things is epistemic assessment if anything is. Consequently, the scale of epistemic positions relevant to their ranking as *epistemically* better or worse is the scale of evaluation which we use when we determine in ordinary life whether people know things. Talk of "epistemic superiority" is thus given its content by *the ideals which govern our ordinary practices of knowledge assessment*. This claim does not by itself beg the question against the view that Infallibility is in fact epistemically superior, since it allows for the possibility that Infallibility is the (unachievable) ideal that governs our ordinary practice. We have to investigate our practices in order to find out whether it is or not. So to determine whether there is an *epistemic* superiority in the position of Infallibility, we must consider the scale of evaluation that we use in ordinary life when we determine whether people know things, and we must ask whether the position of Infallibility appears on that scale (even if we don't ever actually encounter anything which reaches that point on the scale). If it doesn't, then Infallibility

¹¹ See, for instance, p. 307.

isn't *epistemically* superior to the epistemic position involved in knowledge, no matter how desirable it might otherwise be. (Recall again the lesson of the basketball example.)¹²

Two distinct issues must be considered in order to determine the relation of Infallibility to the scale of evaluation that governs ordinary epistemic assessment. First, in our ordinary practices of knowledge assessment, what is required for “eliminating” or “ruling out” an alternative? Second, does the scale of evaluation which governs our ordinary practice always treat it as an improvement if one increases the range of alternatives which one can rule out? Both issues need to be considered, because the Infallibilist view both (1) requires one to be able to rule out *every possible* alternative, and (2) must interpret “ruling out” as requiring possession of reasons or evidence *logically and metaphysically incompatible* with the possibility in question. (A position which satisfied the first requirement, but not the second, would not be an *infallible* position, since even though one could rule out each possible alternative, for any given alternative one's position would be logically and metaphysically compatible with that possibility's obtaining. A similar point applies if only the second requirement is imposed but not the first. For in that case, one's position would be logically and metaphysically incompatible with certain possibilities, but not with others, and so one's position would not be infallible.) Consequently, to determine the relation of the position of Infallibility to the epistemic scale of evaluation, we must determine the relation of these two requirements to the scale of evaluation which governs our ordinary assessments of whether people know things. As I will argue, Infallibility figures rather poorly on each score: the Infallibilist's interpretation of “rule out” does not correspond with what we find in

¹² As this argumentative strategy should make clear, Hetherington is wrong to say that what motivates my view is a commitment to *absolutism*, the view that knowledge does not come in degrees or kinds. One could accept a view such as the one I defend here and *also* accept a gradualist story. What one would have to deny is the idea – in itself inessential to gradualism – that Infallibility is the highest degree on the scale of epistemic evaluation. My fundamental commitment is consequently to the claims (1) that we have no clear idea what's meant by “epistemic superiority” except in relation to our ordinary practices, and (2) that these practices do not treat Infallibility in the requisite way.

My argument for the first claim perhaps warrants further discussion. Here's another way to put it. If talk of ‘epistemic superiority’ is to have any content at all, there must be a scale that is relevant to the ranking of epistemic positions as *epistemically* better or worse. But what scale is that? What epistemic position is at its apex? What is the relation of that position to the position(s) we are content to regard as adequate for knowledge in ordinary life? Saying, “It's the scale of *epistemic* evaluation,” does not provide any guidance for answering these questions. To determine what scale is the relevant one, we have to appeal to two facts. (1) The position we occupy when we have (what we ordinarily call) knowledge lies somewhere on the relevant scale. (2) The kind of assessment we engage in when we carefully and conscientiously consider in ordinary life whether a given person knows a given thing *aims at locating the person's position on the relevant scale*. So, for instance, consider an ordinary sort of case in which we conclude that evidential position E1 is better than E2, but that neither is adequate for knowledge because knowledge requires something better than either. I claim (A) that this sort of assessment is an instance of *epistemic* assessment, and (B) that the scale on which we are placing the positions in this sort of assessment is the scale of *epistemic* superiority. Consequently, what we ordinarily regard as an improvement (or not) when we are being careful and conscientious reveals the terms in which epistemic positions are to be evaluated as epistemically better or worse and the basis for doing so. The ideal(s) which govern our ordinary assessments of whether people know things thus reveal what the ideal epistemic position is.

It might be worried that because our ordinary practices only concern whether people have (ordinary) knowledge and not whether they have attained any *better* position, an investigation of our ordinary practices begs the question at hand; as it might be put, “*Of course* Infallibility will not appear any better than knowledge, in relation to assessments of *that* sort.” However, if we accept that the scale on which epistemic positions are located in evaluations of relative superiority (below the position of knowledge) is the scale of *epistemic superiority*, and that this scale concerns a *uniform dimension*, then our ordinary judgments regarding what does, or does not, constitute improvement on this scale will reveal what the ideal epistemic position is. (I am grateful to Stephen Hetherington for raising a version of the objection addressed here.)

our practice, and an increase in the number or range of alternatives one is able to rule out is not always regarded as an improvement.

Consider, first, what we require in ordinary life for being able to “rule out” an alternative. If one follows David Lewis¹³ in taking a possible not- p world to be eliminated by one’s evidence (memory, sensory experience, etc.) if and only if it is a world in which one does not have that evidence, then being able to rule out every possible alternative to the truth of p would yield infallibility regarding p . For if every not- p world is eliminated by one’s evidence, then there is no not- p world in which one has that evidence. Consequently, the only possible worlds in which one has that evidence are worlds in which p is the case. By contrast, however, if one understands “ruling out” or “eliminating” in such a way that it is a *fallible* relation – so that being in a position to “rule out” or “eliminate” the possibility that q is logically and metaphysically compatible with q ’s obtaining – then being in a position to eliminate or rule out every possible alternative to the truth of p will be a *fallible* position. So even if our ordinary scale of epistemic appraisal is defined in terms of the *range of possible alternatives which one is able to rule out*, the apex of that scale need not be an infallible position. It may be that the highest position on that scale is one which is logically and metaphysically compatible with the falsity of the belief in question.

According to my account, when we take there to be some reason in favor of a certain alternative, we think it inappropriate to attribute knowledge to a person unless he or she possesses decisive positive evidence against that alternative. The phrase ‘decisive positive evidence’ is just a place-holder for the sort of evidence that is accepted in our ordinary practice. For instance, consider the sort of evidence you must have in order to eliminate by observation the possibility that a given bird is of a kind Y which there is some reason to think it might be. Suppose that the question is whether a certain bird is a juvenile cardinal, and that you are attempting to rule out the possibility that it is a house finch. It will suffice if you recognize some feature of the bird which, as things are, is not possessed by house finches, for instance, the presence of a red crest. This kind of evidence is fallible, however, since it is logically and metaphysically compatible with the possibility that the bird is a house finch. Accordingly, as the notion of “ruling out” or “eliminating” figures in our practice, even being able to rule out or eliminate every possible alternative to the truth of one’s belief would not yield infallible belief. I will return to the significance of this point in a moment. (For clarity, from here on I will use ‘rule out_{ordinary}’ as a label for the relation that figures in our everyday practice.)

I turn now to the second question. If the position of Infallibility is the apex of the scale of assessment which we use in ordinary life, then this much at least should be true: the wider the range of alternatives one is able to rule out_{ordinary}, the better one’s epistemic position — as judged from the perspective of our ordinary practices of epistemic appraisal. For if this is not what we judge in our ordinary practice, then our practice is not guided by a scale of assessment of the relevant sort. But in fact, this is not what we judge. We don’t always regard it as an improvement in one’s epistemic position when one increases the number or range of alternatives one is able to rule out_{ordinary}.

To make this point, I want to consider an example. First, however, I should clarify the *point* of appealing to the example. My goal here is not to “pump intuitions” about philosophically contested concepts or theories. Rather, I am trying to make a point about our ordinary practices of epistemic appraisal. Consequently, what I mean to elicit is precisely the judgment that you would actually make regarding this example *in ordinary life*. Our consideration of the example won’t address the issue under discussion unless we react to it in a way that instantiates our ordinary practices of epistemic assessment.

¹³ Elusive Knowledge, *Australasian Journal of Philosophy*, vol. 74 no. 4, 1996, p 553.

Now for the example. You lock your front door, firmly shut it as you walk out, and then carefully check that it is locked, turning the handle and giving the door a good push. Nothing happens. You do it again. Nothing happens. There is nothing funny about your door or lock, and you don't have any reason to think that there is. Your memory is fine; you remember perfectly clearly what has just happened. Would it be any epistemic improvement if now, one moment later, you checked the lock a third time before walking away? Obviously not. There comes a point at which the acquisition of further evidence is no longer epistemic progress. This is such a point — or maybe you have already passed it.

Notice, however, that if you don't check the door a third time, there are possibilities which you are foregoing the opportunity to rule out_{ordinary}. For instance, someone could have secretly replaced your lock with one which silently unlocks after having been checked twice. Checking a third time would have given you decisive positive evidence against this possibility. But acquiring such evidence, we ordinarily judge, is no improvement in the case described. So it isn't always an epistemic improvement to increase the range of alternatives which one is able to rule out_{ordinary}.

There can be some tendency to resist this conclusion. For instance, one might think as follows. "If I were in a world in which locks which silently unlock themselves after a second check were sometimes installed without the owners' knowledge, then checking a third time would be epistemic improvement, since it would enable me to determine that this possibility did not obtain in my case. So here and now, checking my door a third time improves my epistemic position slightly: it provides me with a position which would be secure even if I was in such a world." The trouble with this line of reasoning is that one is not in such a world. And what would be an improvement in one's position if the world were different in certain respects need not be an improvement given the world as it is. If one were in a world in which most basketball players were sixteen feet tall and baskets were set at twenty-five feet, then being eighteen feet tall would be an improvement over being six feet tall for the purposes of basketball. But in the world as it is, being eighteen feet tall would not be an improvement at all.

Still, it's tempting to think that checking a third time is epistemic improvement even in this world, since it inoculates one against the possibility of misleading evidence. If someone seemingly sane and well-informed were to claim that now, after having been checked twice, your door was unlocked, you could reply that you just checked it a third time, and it wasn't. So aren't you better off if you check the door a third time? Well, no. Not if no one seemingly sane and well-informed is around claiming such things. This is not to deny that inoculation against the possibility of certain misleading evidence isn't sometimes an epistemic improvement. It's only to insist that we wouldn't ordinarily judge it an improvement in this case. And that is all that is needed to make my point. (Of course, we would all agree that if there were some evidence — even evidence of which you were not aware — that your door was now unlocked, then checking a third time would be an epistemic improvement. But that is not the case we were initially imagining.)

As I have argued elsewhere, our reaction to examples like this cannot be explained by appealing to purely practical factors, that is, to our interests, purposes, etc., or to those of the person being evaluated, since given our background information about the behavior of door locks, no shift in the purely practical situation would warrant judging it an improvement to go on checking the door.¹⁴ Nor can our reaction to the example plausibly be explained in terms of considerations about purely epistemic costs and benefits.¹⁵ For in the sort of case I had in

¹⁴ This is argued in "Is Infallibility an Epistemological Shortcoming?" pp. 248–9.

¹⁵ That it might be so explained was urged on me by Ram Neta.

mind, there is no epistemic cost to turning the handle again; there is no information of any significance which one thereby misses but might otherwise gain.

My favored explanation of our reaction begins from the observation that in cases of this sort we also judge that there is no reason whatsoever in favor of the alternatives in question; we don't judge it an improvement to seek out evidence against possibilities which we all agree have no reason in their favor. So I suggest that there is no need to rule out_{ordinary} these possibilities precisely because one recognizes that there is no reason whatsoever to think that they obtain. We could perhaps call this, too, a form of ruling out, saying that one can rule out an alternative by recognizing that there are no reasons whatsoever in its favor. But to do so would obscure a difference in kind. What this case has in common with ruling out_{ordinary} is just this: in both cases, the adequacy of one's epistemic position is determined by one's appreciation of the objective situation regarding the reasons that there are for or against the belief in question, and in both cases the position in question is fallible. But the differences are such that it would be extremely misleading to say that the ideal position according to our ordinary standards of epistemic assessment is a position in which one is able to rule out every alternative to the truth of one's belief. It is rather a position in which (1) one is able to rule out every alternative which has some reason in its favor and (2) one recognizes that there is no reason whatsoever in favor of any others.

Here is a way of making sense of our ordinary practice in this regard. (What I offer here requires elaboration and refinement, but it's enough for present purposes.) Think of the reasons that there objectively are for or against the truth of a given proposition p as being determined by what else is the case and what those things tell for or against. (The relevant notion of "telling for or against" here is related to the notion of a *natural indicator* or *sign*.) In the limiting case, as in the case of the possibility that the door will unlock itself after having been checked twice, nothing whatsoever is the case which tells in favor of the possibility's obtaining. In such a case, the objective likelihood of the possibility's obtaining, given everything else that is the case, is extremely slight. Still, a complete specification of everything else that is the case would be logically and metaphysically *compatible* with the possibility's obtaining. Consider, then, an epistemic position (with regard to a particular proposition p) characterized as follows: knowing that nothing is the case which tells in favor of the truth of p . This would be a fallible position, since it is compatible with the truth of p . But if we think of the strength of a person's epistemic position as being determined by the degree to which the position makes it objectively likely that the belief is true, then this position would not be improved if one obtained further evidence against the possibility p . If nothing whatsoever is the case which tells in favor of the truth of p , then in an interconnected, law-governed world such as ours, p will be false. For if p were true, something else would be the case which told in its favor. So if nothing whatsoever is the case which tells in favor of p , then p is not rendered any *less* likely by the specification of some particular fact which tells against it; of course other things are the case which tell against it, given how our world works. That's why it is no improvement in one's epistemic position to obtain evidence against possibilities which are known not to have any reason in their favor.

Even if this explanation is felt to be unsatisfactory, the fact remains that our ordinary practices of epistemic evaluation do not look as they should if Infallibility were a position on the scale of epistemic appraisal. When we decide whether people know things, we consider whether they are able to rule out_{ordinary} those possibilities of error which we take to have some reason in their favor. The rest we ignore because we take there to be no reason in their favor at all. The ideal which governs our practices of epistemic appraisal is thus a position in which one correctly grasps the objective situation regarding the reasons for and against the truth of p , in such a way as to be able to muster decisive but fallible positive

evidence against those alternatives which have some reason in their favor while recognizing that no other alternative has any reason in its favor at all.

This position is obviously quite different from the position of Infallibility. There is no natural way to construct a uniform scale of “improvements” on this position which terminates in the position of Infallibility. And Infallibility is not higher than this position on a scale or ranking of positions which informs our ordinary practice, for Infallibility demands taking account of certain possibilities in a way that would ordinarily be regarded as no improvement at all. Consequently, since the notion of *epistemic* superiority is given its content by the ideals which govern our ordinary practices of epistemic appraisal, we should conclude that however desirable Infallibility might be, it is not *epistemically* superior to the epistemic position which is involved in knowledge. For this reason, Hetherington’s conception of the relation between the Infallibilist skeptic’s claims and our ordinary knowledge ascriptions is deeply mistaken. The Infallibilist skeptic does not reveal that the best we can achieve is epistemically disappointing, for on the scale of epistemic appraisal *nothing* is higher than the ideal which informs our ordinary practice.

There is a perennial philosophical temptation to see our situation in everyday life as both regrettable and blinkered, as if we were prisoners who have grown so accustomed to captivity that we no longer notice it. A central teaching of Austin, Wittgenstein, Quine, and Sellars, as I read them, is that in epistemology at least, this temptation is nothing but a will-o’-the-wisp: when we try to make the charge perfectly clear, it vanishes before our eyes. My goal in this paper has been to develop a detailed version of this general viewpoint. I have tried to show that if we start from the claims and practices to which we are committed – and where else are we to start? – then not only can we reasonably reject Infallibilist skepticism, but we can also see that there is nothing to the charge that because we are not God, even our best is epistemically disappointing.¹⁶

¹⁶ I am grateful to Mark Kaplan and Jonathan Weinberg for helpful conversations and to Ram Neta and Stephen Hetherington for comments on a draft of the paper.