

Contextualism and Fallibilism

Keith DeRose

Contextualism in epistemology has been intimately related to fallibilism, though, due to some malleability in what's meant by "fallibilism," that relation has been presented in different ways. One of contextualism's "founding documents," Stewart Cohen's (1988), presented "fallibilism" as a form of sensibleness in epistemology, and contextualism as a way of achieving it; indeed, the paper was entitled "How to Be a Fallibilist."¹ David Lewis's contextualist manifesto, "Elusive Knowledge" (Lewis 1996), by contrast, memorably construed "fallibilism" as a form of "madness," and presented contextualism as providing a way to "dodge the choice" between it and the even more intrusive madness of skepticism—to steer a course "between the rock of fallibilism and the whirlpool of skepticism" (Lewis 1996: 550). The apparent difference here proves to be merely verbal, based on different uses of "fallibilism" (that we'll be in a position to quickly identify at the end of sect. 2).

In what follows, I will attempt to explain the relations among contextualism, skepticism, and different construals of "fallibilism," and I will distinguish between two contextualist ways of handling what we can call "infallibilist tensions," presenting them along with other, non-contextualist options for dealing with those tensions.

1. Contextualism, Skepticism, and Intuitive Fallibilism

There seems to be an intuitive, but difficult-to-get-precise-about, sense in which we humans are fallible with respect to everything, or at least nearly everything, that we believe, and "fallibilism" is sometimes used to designate this fact about us.² But we will here be interested in uses of "fallibilism" in which it instead asserts that we can *know* things with respect to which we are fallible. "Intuitive fallibilism" can then be the position that knowing some fact is compatible with being fallible with respect to that fact in the murky-but-intuitive sense in question. Following Lewis's lead (Lewis 1996: 449-450), then, we can set aside the issue of whether there are some narrow class of beliefs with respect to which we are infallible—like perhaps a few concerning some simple necessary truths and some truths about our own present conscious experience.³ For, whether or not we are infallible with respect to those special truths,

¹ For a brief history of contextualism, focusing on its early years, and, of course, not at all covering any time after 2009, see (DeRose 2009: 26-29).

² At (DeRose 1990: 289-92), I argue that this fallibility is not just a feature of the *human* condition, but would afflict any cognitive agent with beliefs that we can conceive of. Trent Dougherty (independently) expresses similar thoughts at (Dougherty 2011: 131-2).

³ For the record, I am among those who think that in the intuitive sense in question, we are fallible with respect to all our beliefs, with no exceptions.

common sense demands that our knowledge extends beyond those, to other truths, including many about the external world, with respect to which we seem to be fallible in the intuitive sense in question.

To see contextualism's relation to fallibilism, we look to contextualist treatments of skepticism. On typical contextualist analysis (on the "basic contextualist strategy" I outline at DeRose 1995: 4-7), the skeptic, in presenting her argument, manipulates the semantic standards for knowledge, thereby at least threatening to create a context in which she can truthfully say that we know nothing or very little. The hope is that by means of such an account, our ordinary claims to "know", and ordinary thoughts that we do "know" plenty, can be safeguarded from the apparently powerful attack of the skeptic, while, at the same time, the persuasiveness of skeptical arguments is explained.⁴

The standard contextualist treatment of skepticism then involves the claim/admission that there are standards for knowledge according to which the skeptic is right that we "know" nothing (or perhaps very little).⁵ To employ a handy bit of semi-technical terminology,⁶ the contextualist's is not an account on which we *simply know* what the skeptic denies we know—that is, it does not claim that any speaker using standard English, whatever their context, would be speaking truthfully if they said we "knew" these things. But it is an account on which almost all of our claims to "know" the items in question are true.⁷

This suggests a stance toward intuitive fallibilism. The contextualist (who takes the standard contextualist approach to skepticism) is not what we might call a "simple intuitive fallibilist": they will not hold that all of what gets called "knowledge" in standard English is compatible with being fallible with respect to the beliefs in question. But they will be a "relaxed intuitive fallibilist," holding that our ordinary claims to "knowledge" and ordinary thoughts to the effect that we "know" things (and indeed all such claims and thoughts that are not governed by the peculiar standards of philosophical skeptics) are compatible with our being fallible in the intuitive sense with what we say and think is "known." It also suggests the hope that the sense in which "knowledge" is incompatible with one's being fallible with respect to the item "known" can be used to explain (away) the phenomena that seem to support infallibilism.

⁴ This anti-skeptical story is told much more fully, and defended, in (DeRose ms.).

⁵ Whether this is a claim or an admission will vary from anti-skeptic to anti-skeptic. Some may reluctantly admit it, as the price they must pay for an at least somewhat undermining account of the skeptic's power. Others, of a less strident anti-skeptical orientation, may happily accept it. I tend to fall in the second category.

⁶ Here I use "simply know" as I do at (DeRose 2009: 228-9).

⁷ For discussion of whether contextualist accounts provide sufficient comfort, or whether they leave the most important skeptical worries unaddressed, see especially the "Irrelevant to Traditional Epistemological Reflection on Skepticism?" section of the "Contextualism and Skepticism" chapter of (DeRose ms.).

2. Attempts at Characterizing Intuitive Fallibilism, and the Distinction between Intuitive Fallibilism and GC-Fallibilism

But what is the intuitive way in which we are always, or almost always, fallible with respect to our beliefs? One prominent way of trying to spell this out construes our fallibility as a matter of our always, or almost always, failing to hold our beliefs on the basis of reasons or evidence that entails that those beliefs are true. “Intuitive fallibilism” (in its simple and relaxed varieties) would then be the position that we can know things to be true even when the evidence or reasons on which we base our beliefs in them don’t entail the truth of those beliefs. This is the characterization that Cohen uses; “How to Be a Fallibilist” opens with these words:

The acceptance of fallibilism in epistemology is virtually universal. Any theory of knowledge that endorses the principle that S knows q on the basis of reason r only if r entails q, is doomed to a skeptical conclusion. Fallibilist theories reject this entailment principle thereby avoiding this immediate skeptical result.⁸

Cohen then offers contextualism (and, more particularly, a contextualist version of the relevant alternatives account of knowledge) as a way to handle the problems that confront such a fallibilist. However, I think that characterizations like Cohen’s fail to really capture the intuitive distinction here.

Notice first that “infallibilism,” as Cohen is construing it, does not actually by itself doom us to skepticism. A view that demands a tight connection between our reasons and any knowledge we might have—even one that demands that the tie be maximally tight—does not by itself ensure any strong skeptical results, since by itself it leaves it open that our reasons might abound. In an extreme case, if everything we normally took ourselves to know was among our reasons, then “infallibilism,” construed as we are currently considering it, would not at all threaten *any* of our presumed knowledge. Similar remarks would apply to characterizations phrased in terms of “evidence,” rather than “reasons”—and might be made a bit more urgent by the existence of Williamson’s fairly prominent account of evidence on which $E=K$, as he puts it: our evidence is what we know to be the case (see Williamson 2000: 184-208). Skepticism would result from such forms of “infallibilism” *combined with* some suitably restrictive account of what our reasons or evidence might be.

And this points to one of the problems of this way of distinguishing between fallibilism and infallibilism. In addition to the problems such accounts face in their application to what seems to be our often shaky, fallibilist knowledge of necessary truths,⁹ such accounts have trouble (that can overlap with the just mentioned problem) with *immediate knowledge*. Often, it seems, we know some proposition q on the basis of a reason, or piece of evidence, r, where r is some proposition distinct from q. And

⁸ Cohen 1988: 91. Such a formulation of fallibilism is not new to Cohen. Feldman, for instance, gives such a construal of fallibilism at (Feldman 1981: 266-7). See (Stanley 2005: 127) for an example of such a view which uses “evidence” where Cohen uses “reasons”.

⁹ Baron Reed succinctly sums up this long-noted problem: “[W]here it is necessarily true that p, every justification will entail that p. But this will be so simply because everything entails a necessary truth” (Reed 2012: 586).

perhaps we know r to be the case on the basis of some yet other propositional reason, r_2 . But arguably, as we trace the lineage of reasons back, this basing of knowledge on deeper reasons comes to an end, and arguably, it can end in immediate knowledge, where some subject knows some proposition p to be the case, though her knowledge of p isn't based on any deeper propositional reasons (or any deeper propositional evidence). But it seems that, and nothing clearly rules it out that, immediate knowledge, so understood, can be the kind of shaky, uncertain "knowledge" that an intuitive fallibilist, but not an infallibilist, will accept as such. But how can a formulation of the type we are considering handle such cases? If we say that in cases of immediate knowledge of p , S has *no* propositional reason (or propositional evidence) for p (and so, presumably, doesn't need propositional reasons or evidence for p in order to know that p), then we of course can't differentiate fallibilist from infallibilist immediate knowledge in terms of whether S 's propositional reasons (or evidence) for p entails p . And the other option would seem to be to say that in cases of immediate knowledge of p , p itself is S 's reason (or evidence) for p . But then, of course, p *will* be entailed by S 's reason for it in all cases of immediate knowledge, no matter how shaky they might be. The way of drawing the distinction exemplified here by Cohen would seem tenable only if we assume (rather boldly, I would think—indeed recklessly, relative to my sense of what could be said in defense of such an assumption) that our propositional reasons or evidence (or perhaps our "ultimate" propositional reasons or evidence, if you want to go that way) are themselves always things we know infallibly.

So it is natural to instead try to use a notion like that of epistemic *risk*, or *chances* or *possibilities* of error to draw the distinction: The infallibilist can be construed as holding that, while the fallibilist denies that, in order for S to know that p , there must be no risk (or perhaps "no risk whatsoever"), or perhaps no possibility or no chance ("whatsoever"), from S 's point of view, that p is false.¹⁰ These notions seem capable of applying to immediate as well as mediate knowledge, so as to differentiate in both places the infallibly certain from the at-least-somewhat-shaky.

The problem here is that (at least without the cryptic "whatsoever" qualification), such formulations misclassify some folks, including me, who seem to be intuitive fallibilists, as infallibilists. I take myself to be clearly an intuitive fallibilist, at least of the relaxed variety, but the way of drawing the distinction under consideration would style me an infallibilist. In the case of epistemic possibilities, I am on record as accepting an account on which S 's not knowing that p is a truth condition for an assertion by S of "It's possible that not- p_{ind} ", and, relatedly, as holding that conjunctions of the form "I know that p , but it is possible that not- p_{ind} " express "genuine inconsistencies."¹¹ My contextualism does muddy the waters a bit here. I'm a contextualist about both epistemic modal statements and knowledge attributions, thinking that the meaning of these two types of statements "sway together", in Lewis's nice

¹⁰ In a related move, Trent Dougherty explicates a notion of "epistemic probability," and construes fallibilism as the claim that we can know even where the epistemic probability for us of what is known is less than 1 (Dougherty 2011: 140-42).

¹¹ (DeRose 1991: 596-601). The "ind" subscript indicates that P is to be kept in the indicative mood: "It's possible that she's the best"; not "It's possible that she should have been the best."

phrase,¹² so that, where the epistemic standards are held constant throughout the conjunction, as they are in normal contexts, “I know that p, but it’s possible that not- P_{ind} ” is inconsistent. So I think “I know that p” will often be true even though “It’s possible that not- p_{ind} ” is true as evaluated by *some* epistemic standard—but not the one at which “I know that p” is also true. Still, I seem a good (if muddy) example of someone who thinks that “I know that p” is inconsistent with “It’s possible that not- p_{ind} ,” since I think their relation is much like that of “I am tall” and “I am not tall”—both of which could also be true if evaluated at different standards for tallness, but cannot be if the standards are kept constant.

We might want a way of marking the distinction between those who do and those who don’t think that there is a real conflict between the likes of “I know that p” and “It’s possible that not- p_{ind} ”, and one might well use “infallibilist”/“fallibilist” to mark that distinction, rightly putting me in the “infallibilist” camp. We can call this “GC-”, for “genuine conflict”, fallibilism and infallibilism. But then we should clearly distinguish this GC- fallibilism/infallibilism from the “intuitive” use of the terms. Having made this distinction, we will try to characterize the “intuitive” division in section 3, and will look a bit at the boundary between GC-fallibilism and GC-infallibilism in section 4.

But first: We can now clear up the verbal disagreement between our contextualist authors (noted at the opening). The sensible fallibilism that Cohen seeks is intuitive fallibilism. Contextualists hope that their views can help make the world safe for such intuitive fallibilism—at least in its relaxed variety. The fallibilism that Lewis thinks sounds like “madness” is GC-fallibilism. We will consider how “mad” it is in sect. 4.

3. Characterizing Intuitive Fallibilism in Micro- Terms

Perhaps we can characterize intuitive fallibilism as the view that allows (simply, or in our more relaxed way) that we can have knowledge despite the existence of possibilities or “*micro-possibilities*” (or “micro-risks” and “micro-chances”) of error, in the way I’m about to explain. This attempt certainly won’t render our distinction maximally clear, but it can serve to forge a reasonable account of the relation between the intuitive notion of fallibility and some terms that employ an important part of our epistemic vocabulary (possibility, risk, chance, in epistemic uses).

Consider Unger’s old comparison of “knows” with “flat.”¹³ In the following section, we will consider some tensions between claiming to “know” something and admitting that there is some “possibility”, “chance”, or “risk” of error about the matter, from the speaker’s point of view. Unger

¹² Lewis, writing about the link between similarity and counterfactual conditionals: “I am not one of those philosophers who seek to rest fixed distinctions upon a foundation quite incapable of supporting them. I rather seek to rest an unfixed distinction upon a swaying foundation, claiming that the two sway together rather than independently” (Lewis 1973: 92). See (DeRose 2009: 19-20) for a little discussion of applications to epistemological language.

¹³ (Unger 1971); updated in Chapter 2 of (Unger 1975).

noted similar tensions between calling a surface “flat” and admitting that it had “bumps”. Largely from such tensions, he concluded if a surface is flat, then it has no bumps whatsoever. However, Unger claimed, just about all the physical surfaces we encounter, including those we typically describe as “flat”, do in fact have some bumps, however small (and sometimes it might take a microscope to see them), and so, he argued, they are not actually flat (Unger 1975: 65). Unger concluded that we are just about always speaking falsely when we call an ordinary physical surface “flat”.

Those who cling to the truth of our ordinary claims that surfaces are “flat” have several options as to what to say about Unger’s argument and the microscopic “smallish bumps” he points out. Employing our semi-technical use of “simply”, I suppose that one could bravely declare that the “irregularities” revealed by microscopes *simply* are not bumps, and anyone who calls them “bumps” in any context is speaking falsely. Those things simply aren’t big enough to really constitute “bumps”! Alternatively, one could agree with Unger that it turns out that almost all the physical surfaces we talk about, including those we typically call “flat”, simply *do* have bumps, but hold that being flat, and being truthfully called “flat”, is consistent with having (small enough) bumps. But many will be tempted by a contextualist approach to “bumps”, saying that what counts as a “bump” varies with context. In most contexts, the microscopically small “irregularities” Unger points to don’t rise to the level of counting as “bumps”—which is why we do (despite our standing knowledge of what the surfaces of physical objects are like) and can truthfully say of many ordinary physical surfaces in many contexts that they have “no bumps”. (Or even, in some perfectly natural sense, “No bumps whatsoever”: You worked for a long time sanding a tabletop to remove all the large bumps, but still left about ten “small bumps,” as you were happy to call them. The next day, you sand those down. “Now there are no bumps whatsoever,” you say, apparently [or at least this seems apparent to me] truthfully, despite the continued presence of Unger-bumps.) Relative to those contexts, the small “irregularities” that Unger writes of are what we might well call “micro-bumps”: they are things that don’t count as “bumps” relative to the context under discussion, but do count as “bumps” in other contexts, where they meet the lower thresholds in place for what is to count as a “bump”. Such a contextualist line would seem to hold promise for accounting for the persuasiveness (such as it is, and this does seem to vary much from person to person) of Unger-like arguments for flatness skepticism. Perhaps pointing out such small irregularities and calling them “bumps” is a way of putting into place, or at least of pushing toward putting into place, standards by which they do count as “bumps”—of making what were just micro-bumps now count as “bumps”? And here one can sense possibilities (that we won’t here explore) for explaining the tensions that Unger exploits in ways that won’t involve us in any jarringly incredible claims to the effect that almost all our positive ordinary claims to the effect that ordinary physical surfaces are “flat” are false.

It’s easiest to make sense of a notion of “micro-bumps” when one accepts contextualism about “bump”, so one can say, as I do above, that a micro-bump relative to context *c* is something that does not count as a “bump” in *c*, but does count as a “bump” in other contexts with more liberal standards. But suppose that for some reason (perhaps an irrational aversion to context-sensitivity) one is an invariantist about “bump”, and so holds that, though the standards for what we are likely to *call* a “bump” may well vary a lot among contexts, what can be *truthfully* called a “bump” does not vary from context to context. Well, then, you won’t think there are any micro-bumps relative to any contexts—at

least given our current understanding of “micro-bump”. But it seems you might still have good use for a related notion, which might well be given the same label. For, supposing you’re not a *skeptical* invariantist like Unger, but instead think that the standards for what counts as a “bump” hold steady at some moderate level (at which microscopic “irregularities” certainly don’t count), you will often face things much like those you would count as “bumps” but which you’ll think don’t—and that sometimes don’t *quite*—rise to the level of counting as a “bump”, and you might have use for a term like “micro-bump” to describe such things, where by this you will mean roughly: something much like a bump, but which doesn’t rise to the level of being a bump; something that, if there were only more of it, would be a bump.¹⁴ If we can make sense of such a notion (and I admit, it’s not the clearest in the world), that should become our general notion, available to both contextualists and invariantists about “bump”. We can use the explication just given, but relativize it to contexts: A micro-bump, relative to context *c*, is something which is much like the things that count as “bumps” in *c*, but which doesn’t rise to the level of counting as a “bump” in *c*; it is something that, if there were only more of it, would count as a “bump” in *c*. For invariantists, the line between “bumps” and micro-bumps does not move from context to context. For contextualists, it does. This gives the contextualist an added aid in answering questions of the likes of “Wha’dya mean, ‘something that, if there were only more of it, would count as a bump in *c*?’!”: for the contextualist can often add that it is something that *does* count as a “bump” in contexts other than *c*, which are more liberal than *c* in counting things as “bumps”.

Though it’s not a physical notion of something we can visually imagine, we can utilize an analogous notion of “micro-risks” of error, understanding them to be micro-bumps in the epistemic road. A “micro-risk” of error, relative to context *c*, is something which is much like the things that count as “risks” of error in *c*, but which doesn’t rise to the level of counting as a “risk” of error in *c*. It is something that, if there were only more of it, would count as a “risk” of error in *c*. This notion and/or closely related notions may be helpful in characterizing what “infallibilism” about knowledge is: Perhaps we can say that the infallibilist, but not the fallibilist, holds that knowing that *p* is incompatible with there being risks *or even micro-risks* (and/or micro-possibilities, and/or micro-chances) of error with respect to *p*? And perhaps our fallibilist will be able to explain the tensions we looked at in the previous section in ways that won’t involve us in any incredible claims to the effect that even the tiniest micro-risks of error are enough to make even ordinary claims to “know” something go false?

¹⁴ So micro-bumps will often “behave” in ways similar to bumps. So, to use an Unger-inspired illustration, just as you should be mindful of the potential effects of bumps in the field you’re playing on when planning a shot in croquet, or when you’re predicting how likely your planned shot is to succeed, so, especially when it’s an important shot and one that calls for much precision, you might do well to consider what some prominent micro-bumps might do to your shot.

4. Contextualism and GC-Fallibilism/GC-Infallibilism: Options for Handling the Infallibilist's Tensions

We close by looking at the options by which intuitive fallibilists might seek to deal with the “infallibilist’s tensions” (so-called because it is the infallibilist who is likely to point them out): tensions between claiming to know that *p* while, at the same time, making some admission about one’s limited position with respect to a matter (like admitting that there is a risk that you might be wrong about *p*—or that there is a chance or a possibility that *p* is false, or that one may or might be mistaken). Two of the interesting intermediate options here are contextualist in nature (though the other options are also open to contextualists).

Let’s start with a look at the kinds of tensions involved. Here is Rogers Albritton having some fun with how hard it is to combine “I know” with “I may be mistaken”; he has just rejected the claim that “I may be mistaken” requires for its truth that the speaker have what J.L. Austin calls, not just a reason for thinking she may be mistaken, but a “*concrete reason*” for supposing so:

But perhaps the point is that one gives people to understand by saying the sentence [“I may be mistaken”], and leaving it at that, that one does have some concrete reason to suppose that one may be mistaken or wrong, in the case at hand. If so, that “implicature,” in Grice’s term, should be easy to cancel. Like this: “I know he is honest. But I may be mistaken. I have no concrete reason, in this case, to suppose that I may be mistaken, much less that I am mistaken. But then, I didn’t in that other case, either, though as you will recall I was disastrously mistaken, there. So if I were you, I wouldn’t count on it that I’m not mistaken again. Obviously, I may be. I wouldn’t say, “may well be.” This chap is extremely convincing. But he may be dishonest, of course. Nevertheless, as I was saying, I know he’s honest. That’s the position.”

But this “position” doesn’t exist. No implicature, if that’s what it is, has been cancelled. On the contrary, the speaker has absurdly undermined what would have been a pretension to know, if he had said as much and shut up. Or does he know? Perhaps he does, and should have said so at the end, if not at the beginning, more emphatically, in which case he would have cancelled his rambling concession that he might be mistaken.

The fact is, “I know” and “I may be mistaken” can’t be gotten through a logical intersection by adroit steering and some sounding of horns. They inexorably collide.¹⁵

Similar observations would be plausible about the relation of “I know” to “There is some risk that I am wrong”, or “There is some chance that I am wrong”, or to any number of other ways of admitting some

¹⁵ Albritton 2011: 5-6. These remarks are from a talk Albritton gave, I believe at the University of California, Irvine, I believe in 1987, a slightly fixed-up text of which was later published posthumously as (Albritton 2011). Albritton (who was my adviser) viewed the choice before us (in the terms we are here using) as being between the kind of implicature-based GC-fallibilism that he mentions and a quite thorough-going GC-infallibilism. A main push toward the birth of my contextualist views was the pressing sense that there were other—and better—important options here.

kind of possibility of error—or, again, of fallibility, in some good sense, with respect to the thing putatively known.

Of course, some may wonder just how inexorable, and how solid, is the collision of “I know” with “I may be mistaken” and with similar admissions of the chance or risk or possibility of error. My purpose here is to quickly survey the basic positions on such conflicts open to an intuitive fallibilist (of the simple or relaxed variety). Which strategy is most appealing to the intuitive fallibilist in a given case may depend on just how much of a tension you think there is, and how inexorable you find it to be. We will arrange the options according to how fallibilist, or how clearly fallibilist, they are in the GC-sense of that term. We are then arranging them by how genuine these positions take the conflict between the admission in question (e.g., “It’s possible that not- p_{ind} ”) and the claim to know (“I know that p ”) to be. Alternatively, we can think of these strategies in terms of what they say about the skeptical inference from the admission in question to the denial of knowledge, e.g., “It’s possible that not- p_{ind} ; so, I don’t know that p ”.

The reaction to the alleged tensions that would be the most zealously fallibilist in the GC-sense would be the way of flat denial (if I knew the Latin, I could get a zippy label for this way by filling in the blank of “modus ____ tollens”¹⁶): One who simply denied that there is any serious tension between the admission and the claim to know, nor even any appearance of validity to the skeptical inference, to be accounted for would certainly be a GC-fallibilist of the first order. I don’t find that strategy credible in the cases before us (though I have witnessed the stance being (incredibly) adopted), but it should be included in our brief survey of options, not just for the sake of others who might react differently, but for the possibility of extending the scope of these strategy types to cover other admissions, where this strategy might make more sense.

A second stance that is a little less stringently GC-fallibilist, and that can be seen as roughly Moorean in methodological character, would be to admit that there is some slight intuitive pull toward finding the admission incompatible with the claim to know (and toward finding the inference valid), but to hold that it is so dominated an opposing intuitive push in the other direction that there is no call for trying to account for the slight pulls before reaching a secure verdict against them.

On the other extreme, a GC-*in*fallibilist of the first order holds that no standard use of “I know that p ” can be true if “It’s possible that not- p_{ind} ” is true for the speaker at any epistemic standard by which that modal claim can be governed. Though it is in principle possible for a contextualist to hold this view, this kind of GC-*in*fallibilist will likely be an invariantist, holding that for a given subject in a given situation, there is just one epistemic standard that can govern both her knowledge claims and her epistemic modal statements,¹⁷ and that the terms are connected in such a way that one cannot count as

¹⁶ Trent Dougherty now informs me that he majored in Latin and he assures me that the term I'm looking for here is “modus frickin tollens.”

¹⁷ For “Classical Invariantists,” there will be only one set of epistemic standards that can ever govern knowledge ascriptions and epistemic modal statements. For “Subject-Sensitive Invariantists,” we are instead speaking of the only epistemic standards that can be applied to particular subjects given certain features of those subjects’ situations. For the distinction between the types of invariantists, see (DeRose 2009: 23-26).

knowing that p if, according to the epistemic standard that must govern one's relevant modal claims, it is epistemically possible from one's point of view that p is false. Early Unger is a good example of an extreme GC-infallibilist, but one does not have to be a *skeptical* invariantist, nor an infallibilist in the intuitive sense, like him to be a fervent GC-infallibilist. In fact, I suspect this will be where many invariantists will land. A moderate invariantist, perhaps impressed by the apparent sharpness and inexorability of the collisions that occur in our "logical intersections," can adopt this view. Being moderate, she will think we often enough do simply know all manner of things, and will in those cases of knowledge conclude that "It is possible that not- p_{ind} " is simply false when said by one who possesses such knowledge of p. Early Unger combines infallibilism in the intuitive sense with extreme GC-infallibilism; our moderate invariantist, though, shows how such extreme GC-infallibilism can be combined with fallibilism in the intuitive sense.

That leaves the intermediate views, of which I will distinguish four. A view that acknowledges the existence of an at-least-apparent conflict that needs to be dealt with, and so isn't dismissive in the way that the GC-fallibilist views we've so far are, but which claims that the conflict is due to a Gricean conversational implicature generated by either the knowledge claim or by the relevant admission of fallibility, should clearly be classified as a GC-fallibilist view. This is the type of view Albritton considers, but then rejects, in the quotation at the beginning of this section, and has been the focus of some of the wrangling over what we are here calling GC-fallibilism, since Patrick Rysiew presented such an account (Rysiew 2001: 492-8), setting off a battle over the tenability of such a mere "pragmatic" account of (some of) our clashes.¹⁸ A related view would have it that the conflict is generated by a *conventional* implicature of one or both of the claims. This option would seem a bit more GC-infallibilist than the previous one, since on it the conflict is generated by the conventional meaning of the two sentences, even if not by their truth-conditional content. I am a bit leery of conventional implicatures myself (DeRose 2009: 88-9, n. 9), and should perhaps leave the classification of this position on our GC-scale to those who feel more at home with them, but I'm guessing that one who went for such an account of a tension should be construed as a GC-fallibilist.

The two intermediate positions that are left are both contextualist options. As I indicated in the previous section, my own position, at least with respect to the knowledge-epistemic possibility clash, is that the meaning of the two claims is such that what's expressed by "I know that p" is incompatible with what one would express by "It's possible that not- p_{ind} ", *when both are evaluated at the same epistemic standard*. On this view, the inference "It's possible that not- p_{ind} ; so, I don't know that p" is what we can call "Stine-valid": The conclusion does follow from the premise, so long as one evaluates both at the same epistemic standards, and thereby avoids committing, in Gail Stine's memorable words, "some logical sin akin to equivocation" (Stine 1976: 256). As I indicated in section 2, I think this view should be

¹⁸ See Hawthorne (2004: 24-28) and (Stanley 2005) for criticisms of pragmatic accounts of the clash (aimed explicitly at Rysiew only in the case of Stanley), and (Dougherty and Rysiew 2009) for replies. This battle over the tenability of "pragmatic" accounts of our clashes is closely related to my wranglings with Rysiew over the tenability of his "pragmatic" account of the conversational data that is used in support of contextualism, which is the primary focus of (Rysiew 2001); see (DeRose 2009: 118-124).

classified as a GC-infallibilist one, for, after all, on it, “I know that p” seems to be inconsistent with “It’s possible that not-p_{ind}” in as strong a way as “I am tall” is inconsistent with “I am not tall”.

Note, however, that the above option is only available where one is contextualist about both the knowledge claim and the admission of fallibility, so that the meanings of the two sentences can “sway together”. There is an importantly different contextualist position available even for (but not only for) cases where the knowledge claim is subject to varying epistemic standards, but the admission of fallibility is not—or *vice versa*. On this position, the inference from the admission of fallibility to the admission that one does not know is not Stine-valid, but rather what we can call Stalnaker-reasonable. To illustrate, suppose you are an invariantist about, say, “There is at least some slight chance that not-p”, thinking that such a claim is not governed by varying epistemic standards, and, say, you think it is almost always true. Being a good fallibilist in the intuitive sense, you’ll think that people often speak truthfully when they claim to know things, even where, as is almost always the case (on your view), there is at least some slight chance from their point of view that the things they are claiming to know are false. You can then hold that “There is at least some slight chance that not-p; so, I don’t know that p”, while invalid, is a “reasonable inference” in something like the sense that Robert Stalnaker proposed in his (1975). To say that inference is “Stalnaker-reasonable” is to say that the assertion of the premise affects the meaning of the conclusion so that the conclusion will, if need be, (tend to) come to express a proposition that must be true if the premise is true. In the case of our sample inference, you might think that bringing up the matter of a slight chance that not-p will be most germane to, and for that and/or other reasons may invoke, standards for knowledge at which slight chances of not-p are enough to block “knowledge” of p. On this view, then our sample skeptical inference is not valid: The conclusion can be false where the premise is true—so long as one keeps quiet about that true premise. Rather, the assertion of the premise affects (or at least has a tendency to affect) the meaning of the conclusion in such a way that the latter comes to follow from the former.¹⁹ Moving from the inference to the corresponding tension, on this type of view, putting “There is at least some slight chance that not-p” into play tends to create (if need be) a context where “I know that p” is governed by standards that render it incompatible with (even) slight chances that not-p. In light of its ruling that knowledge can co-exist with the truth of the relevant claim of fallibility so long as that latter isn’t actually made, I would classify this view as a GC-fallibilist one. That the knowledge-asserting sentence comes to (or has a tendency to come to) express a proposition incompatible with the admission of fallibility when the latter is put into play would be best thought of as an account of why knowledge of facts can misleadingly

¹⁹ In his (2015), Alex Worsnip provides a great example of a GC-fallibilist view of this type, but one on which it is the knowledge claim’s effect on the meaning of the epistemic modal statement, rather than the other way around, that accounts for the clash in trying to make the claims together. I think Worsnip’s should be classified as a GC-fallibilist view, because, though he thinks that “I know that p, but it’s possible that not-p” is inconsistent *wherever it is asserted*, on his view, the knowledge claim and the epistemic modal claim, and indeed their conjunction, can all be true together relative to many contexts (in which these claims are not uttered). That they cannot be truthfully asserted together accounts for why they misleadingly appear to be inconsistent. This grounds an important sense on which, for Worsnip, as opposed to a GC-infallibilist like me, there can be “Possibly False Knowledge.” I don’t accept Worsnip’s account because I don’t find his (4) – (7) (at 2015: 232) as felicitous as he does.

appear to be incompatible with, say, even slight chances to the contrary. This form of GC-fallibilism may be wrong (and I think it is), but it seems far from “madness.”

The line between GC-fallibilism and GC-infallibilism, then, seems to me best placed so that it cuts right between our two intermediate, contextualist options. But we can all be good fallibilists in at least the relaxed version of the intuitive sense of that term.²⁰

References

Albritton, Rogers 2011. “On a Form of Skeptical Argument from Possibility,” *Philosophical Issues* 21: 1-24.

Cohen, Stewart 1988. “How to be a Fallibilist”, *Philosophical Perspectives* 2: 91-123.

DeRose, Keith 1990. *Knowledge, Epistemic Possibility, and Scepticism*, Ph.D.diss., University of California, Los Angeles; University Microfilms International.

_____ 1991. “Epistemic Possibilities,” *Philosophical Review* 100: 581-605.

_____ 1995. “Solving the Skeptical Problem,” *Philosophical Review* 104: 1-52.

_____ 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1*, Oxford: Oxford University Press.

_____ ms. *The Appearance of Ignorance: Knowledge, Skepticism, and Context, Vol. 2*.

²⁰ Thanks to Patrick Rysiew for helpful comments on a draft of this article, and especially to Trent Dougherty for comments and much very helpful e-discussion of these matters.

- Dougherty, Trent 2011. "Fallibilism," in S. Bernecker, D. Pritchard, ed., *The Routledge Companion to Epistemology*.
- Dougherty, Trent and Rysiew, Patrick 2009. "Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions," *Philosophy and Phenomenological Research* 78: 123-132.
- Feldman, Richard 1981. "Fallibilism and Knowing that One Knows," *Philosophical Review* 90: 266-82.
- Hawthorne, John 2004. *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Lewis, David 1996. "Elusive Knowledge", *Australasian Journal of Philosophy* 74: 549-567.
- Reed, Baron 2012. "Fallibilism," *Philosophy Compass* 7/9: 585-596.
- Rysiew, Patrick 2001. "The Context-Sensitivity of Knowledge Attributions," *Noûs* 35: 477-514.
- Stalnaker, Robert 1975. "Indicative Conditionals," *Philosophia* 5: 269-286.
- Stanley, Jason 2005. "Fallibilism and Concessive Knowledge Attributions," *Analysis* 65: 126-131.
- Stine, Gail 1976. "Skepticism, Relevant Alternatives, and Deductive Closure," *Philosophical Studies* 29: 249-261.
- Unger, Peter 1971. 'A Defense of Skepticism', *Philosophical Review* 80: 198-219.
- ____ 1975. *Ignorance: A Case for Scepticism*, Oxford: Oxford University Press.
- Williamson, Timothy 2000. *Knowledge and its Limits*, Oxford: Oxford University Press.
- Worsnip, Alex 2015. "Possibly False Knowledge," *Journal of Philosophy* 112: 225-246.