

Appendix E: Attempts to Explicate Intuitive Fallibilism and the Distinction between Intuitive and Genuine Conflict (GC-) Fallibilism

In section 24 of Chapter 5, I suggest that notions like those of micro-risks, micro-chances, and micro-possibilities of error might be helpful in characterizing what ‘fallibilism’ and ‘infallibilism’ about knowledge are, since they allow us to say that the infallibilist holds, while the fallibilist denies, 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*. Here, I will provide a little background as to why one might be driven to such a characterization, and I will also, both here and in Appendix E, try to clarify some of the relations between contextualism and fallibilism.

Contextualism, Skepticism, and Intuitive Fallibilism

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.’ 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’.

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, in what can be called

a ‘condition’ way, to designate this fact about us.¹ But we will here be interested in what we can call ‘compatibility’ 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.² For, whether or not we are infallible with respect to those special truths, 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. As we have seen in this volume, the standard contextualist treatment of skepticism (one that follow the ‘basic contextualist strategy’ I explain in section 2 of SSP) 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). To again use our handy bit of semi-technical terminology, 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.

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

¹ 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).

² 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.

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. And we will look at basic strategies for doing that in Appendix E.

The Non-Entailing Reasons/Evidence Account of Intuitive 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. (Cohen 1988: 91)

Characterizations of these positions that are roughly like Cohen's are initially attractive.³ To give just one other example of the type of characterization I have in mind, Jason Stanley uses the notion of evidence in much the way that Cohen uses reasons, when he writes:

Fallibilism is a certain claim about the character of one's evidence for one's knowledge. Fallibilism is the doctrine that someone can know that p, even though their evidence for p is logically consistent with the truth of not-p. (Stanley 2005b: 127)

However, I think that characterizations like this, on which the dispute concerns whether knowledge requires a logically airtight connection between what

³ Such a formulation of fallibilism is not new to Cohen. Feldman, for instance, gives such a construal of fallibilism at (Feldman 1981: 266-7).

is known and something like the putative knower's evidence or reasons for her claim, fail to really capture the intuitive distinction here.

Notice 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, like Cohen-style 'infallibilism,' 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 it is here, would not threaten *any* of our presumed knowledge. Similar remarks would apply to Stanley's characterization, if we just substitute 'evidence' for '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 (Williamson 2000a: 184-208). Skepticism would result from such forms of 'infallibilism' *combined with* some 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,⁴ such accounts have trouble (that can overlap with the preceding 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 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

⁴ Baron Reed succinctly sums up this 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).

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.

Accounting for Intuitive Fallibilism in Terms of Risk, Chances, or Possibilities of Error?

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.⁵ 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.⁶

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

⁵ 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).

⁶ It can still be held that *in cases of non-immediate knowledge*, a lack of entailment between a believer's reasons (or evidence) for p and p itself is sufficient to render the believer fallible with respect to 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'.⁷ 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 phrase,⁸ so that in any normal context, where the epistemic standards are held constant throughout the conjunction, '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.

Genuine Conflict (GC-) Fallibilism Distinguished from Intuitive Fallibilism

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. For the GC-construal of the distinction does not seem to capture the intuitive distinction we have been seeking to characterize, since it seems that I, and those with views similar to mine, am a GC-infallibilist, while I am not an infallibilist in the intuitive sense, since I think that knowledge is—or at least that what we rightly call 'knowledge' at least almost always is—compatible with the believer being, in the elusive sense we are seeking to characterize, fallible with respect to her belief.

We can now clear up the verbal disagreement between our contextualist authors, noted in the second paragraph of this Appendix. The sensible fallibilism

⁷ (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.'

⁸ 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.

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 Appendix E.

Characterizing Intuitive Fallibilism in Micro- Terms

But this all leaves us still looking for a way of characterizing Intuitive Fallibilism. And it’s here that the notions of ‘micro-possibilities’, ‘micro-risks’, and ‘micro-chances’, as those are used in section 24 of Chapter 5, would seem to be the obvious way of drawing the distinction, in light of the problems we have just encountered. Stated in a way available to both invariantists and contextualists about knowledge attributions, and about the other notions involved, the ‘infallibilist’ will be one who holds that ‘S knows that p’ can never be true if there is, from S’s point of view, even a micro-possibility (or micro-risk or micro-chance) that not-p (where by ‘even a micro-possibility that not-p’ we mean ‘either a possibility or a micro-possibility’, and similarly for ‘even a micro-risk’ and ‘even a micro-chance’). This characterization certainly doesn’t render our distinction maximally clear, but it serves to forge a reasonable connection between the intuitive notion of fallibility and some terms that employ an important part of our epistemic vocabulary (possibility, risk, chance, in epistemic uses), and it seems to capture the intuitive distinction quite well. And indeed, something like this may well be what people are sometimes trying to reach for when they portray the infallibilist as holding that knowledge is incompatible with ‘any risk at all’, or ‘any possibility whatsoever’, or the like, of error.

Note that while I was no doubt moved to draw a distinction between GC- and Intuitive Fallibilism by trying to fit contextualist views such as my own into the intuitive fallibilist camp, this way of drawing the distinction can also be helpful to certain moderate invariantists. Such a character may well, like me, acknowledge a genuine conflict in the likes of ‘I know that p, but it’s possible that not-p_{ind}’, and so, like me, will be a GC-infallibilist. She might in fact have a view of the conflict *very* much like mine, only she will see the epistemic standards involved as ‘staying put together’, as it were, where I see them as ‘swaying together’. In other words, she might accept the same kind of linkage I do between knowledge and epistemic possibility, but, not thinking there’s any swaying of epistemic standards going on

here, she will have no use for the Lewisian thought that the two ‘sway together’. But, like me, she might feel that in the important intuitive sense we are seeking, she should be counted as a ‘fallibilist’. And that she recognizes that there can be knowledge even where there are micro-possibilities, micro-risks, and/or micro-chances of error might for her do a good job of capturing the intuitive way in which she is a ‘fallibilist’. For her, these micro-possibilities of error will not be things that can be correctly called epistemic ‘possibilities’ of error in some extreme contexts. Still, for all that, they can for her be things that are much like ‘possibilities’ of error, but which don’t (and in some case, don’t *quite*) rise to the level of being possibilities (risks, chances) of error.

Note also that, although infallibilism in the sense that is opposed to Intuitive Fallibilism does seem to lead fairly directly to a quite extreme skepticism, it does at the same time seem to have some attractions—which attractions Unger-like skeptics exploit. One such attraction could be the infallibilist feeling that anything that deserves the label of ‘knowledge’ should be something solid enough that you can knowledgeably expand it via deduction, without having to worry about things like the shakiness of your known premises messing up the whole operation. In other words, there can be some attraction to what in Chapter 5 I have called ‘Oxford closure’—even as I have claimed the stronger intuitive force is aligned against it. The characterization of the intuitive distinction that we are now considering seems perfectly suited to account for that attraction of infallibilism, for it is freedom from even micro-risks of error that is needed to underwrite multi-premise deductions as exceptionlessly expanding our knowledge to the conclusions of competent deductions, without regard to the problem of the aggregation of risk.