KNOWLEDGE A Human Interest Story

BRIAN WEATHERSON



https://www.openbookpublishers.com

©2024 Brian Weatherson



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0). This license allows you to share, copy, distribute and transmit the text; to adapt the text for non-commercial purposes of the text providing attribution is made to the author (but not in any way that suggests that they endorse you or your use of the work). Attribution should include the following information:

Brian Weatherson, *Knowledge: A Human Interest Story*. Cambridge, UK: Open Book Publishers, 2024, https://doi.org/10.11647/OBP.0425

Further details about the CC BY-NC license are available at http://creativecommons.org/licenses/by-nc/4.0/

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at https://archive.org/web

Any digital material and resources associated with this volume will be available at https://doi.org/10.11647/OBP.0425#resources

ISBN Paperback: 978-1-80511-394-2 ISBN Hardback: 978-1-80511-395-9 ISBN Digital (PDF): 978-1-80511-396-6 ISBN Digital eBook (EPUB): 978-1-80511-397-3 ISBN HTML: 978-1-80511-398-0

DOI: 10.11647/OBP.0425

Cover image: Golden Gate Bridge, San Francisco, California. Photo by Tj Kolesnik, https://unsplash.com/photos/silhouette-photography-of-mountain-beside-suspensionbridgeduring-golden-hour-Vnz7o7LeuDs.

Cover design: Jeevanjot Kaur Nagpal

The next three chapters are primarily defensive; they are responding to the three objections to IRT that seem to me most serious. But they aren't just defensive. I'm not just saying why the theory from the chapters to date is immune to these arguments. I'm also developing the theory. That's especially true in this chapter, which is why it is first. So what are these objections?

The first is what I'll call the *objection from double checking*. As Jessica Brown (2008) argued, there are plenty of cases where intuitively a person knows that *p*, but should check whether *p* is true. This seems to be a problem for IRT, since it is motivated by the thought that what's known is an appropriate starting point in inquiry. At first glance, it's very weird to have an inquiry into *p*, when the inquirer is in a position to simply say *p*, therefore *p*. I used to think that in these cases the defender of IRT would have to either say that they are not really cases of knowledge, or not really cases of appropriate inquiry. Unfortunately, neither of these options was particularly successful. I now think the objection should be addressed head on. It is possible to properly conduct an inquiry into p, even when one knows that p, and even when knowledge provides appropriate starting points for inquiry. That's because it is often appropriate to deliberately restrict oneself in inquiry, and use fewer resources than are otherwise available. The aim of this chapter is to defend the claims made in the last two sentences, and to show how they provide a response to the objection from double checking.

The second is what I'll call the *objection from close calls*. As Alex Zweber (2016) and, separately, Charity Anderson and John Hawthorne (2019a) showed, some simple versions of IRT say implausible things about cases where a person is choosing between very similar options. What I'm going to argue is that the problem their cases raise is not due to IRT, which is correct, but to the background assumption that choosers

should maximise expected utility. My response is going to be that in the cases they describe, choosers should not maximise expected utility. That might sound like an absurdly radical view, since expected utility theory is at the heart of all contemporary decision theory. But expected utility theory has fairly implausible things to say about *close call* cases. A better theory, one that takes account of deliberation costs, is both more plausible, and consistent with IRT. I'll say much more about this in Chapter 6.

The third is what I'll call the *objection from abominable conjunctions*. This is the IRT-equivalent of the blank stare objection to modal realism. Many people find it simply implausible that knowledge could depend on something like interests, which are not relevant to the truth of what is purportedly known. The defender of IRT owes a reply to this widespread feeling. Part of my reply came back in Chapter 1. I think this feeling is a result of being in a very strange place in the history of epistemology, where the focus is on fallibilist, interest-invariant, concepts. But we can do better than that. It is hard to articulate the intuition behind the unhappiness with IRT without lapsing into the Justified True Belief (JTB) theory of knowledge. Most plausible solutions to the problems with the JTB theory end up introducing kinds of interest-relativity for independent reasons. I'll go over these responses in Chapter 7.

So those are the three objections I'm going to spend a lot of time on. There are three other classes of objection I'm not going to spend much time on.

The first class are objections to IRT that assume that knowledge changes when and only when one is in a 'high-stakes' situation. Since I don't assume that, those objections don't raise problems for my version of IRT.

The second class are objections to IRT that assume that some parts of epistemology are interest-invariant, while some are interest-relative. While I used to endorse such a theory, I no longer do. This book defends a global interest-relativism where knowledge, belief, rationality, and evidence are all interest-relative (in different ways). So these objections don't raise problems for my version of IRT either.

The third class are objections to IRT that only apply to versions of IRT that add on an opposition to contextualism or relativism. With this addition, IRT becomes what has been called *interest-relative invariantism*, or IRI. While I've defended that in the past, I'm not going to defend it

here. The thesis of this book is that knowledge is interest-relative. If you want to understand the word 'knowledge' in the previous sentence in a contextualist or relativist way, go right ahead. Whatever metasemantic theory you have about the kind of words 'knows' and 'knowledge' are, I will be willing to defend the claim that knowledge is interest-relative.

5.1 Starting and Settling

At the heart of the influential picture of inquiry developed by Jane Friedman (2017, 2019b, 2019a, 2020, 2024b) is the view that humans are capable of a number of distinctive attitudes. To be inquiring into some question, she argues, is to have a *questioning attitude* towards that question. That's to say, she does not identify inquiry with particular actions, or at least with particular bodily movements. An actor might mimic the movements an inquirer makes without actually inquiring; a genuine inquirer might be sitting in an armchair quietly synthesising their evidence. So particular movements are neither sufficient nor necessary for real inquiry. Rather, inquiry is a state of mind, a questioning state of mind.

The contrast to having a questioning attitude is having a settled attitude.¹Friedman holds that to believe something is to treat the question of whether it is true as being affirmatively settled, and I'm adopting the same position here. This attitude is deeply related to inquiry. Typically things are settled as the result of inquiry. Also typically, one does not inquire into something one has settled. Friedman makes a further claim: if one does inquire into something one has settled, this is a kind of mistake. It is incoherent to have both a questioning and a settled attitude towards the same question. I'm going to disagree with this further claim, while mostly adopting the broad picture she develops.

The main difference between her picture and the picture of inquiry I'm using concerns where beliefs go in inquiry. I think that treating something as settled is most fundamentally about willingness to use it

¹ These are contrasts, but they don't exhaust the space. One might not have an attitude to a question. I'd also say that one might not treat a question as settled while not inquiring into it, because one treats the question as unworthy of effort, or impossible to make progress on. As Friedman (2024a) notes, it gets complicated to say something coherent about these cases while allowing for the possibility of inquiry to be reopened.

as the beginning of a new inquiry. The essential feature of belief is that it starts inquiry, not that it ends inquiry. What makes an attitude a belief is not that inquiry into it is settled, it's that it can be used in the process of settling open questions. I used to think that whether one identified beliefs with settled states, or with the inputs to inquiry, was only a difference of emphasis, and a pretty minor one at that. After all, beliefs are typically the outputs of one inquiry and then serve as inputs to another; whether one takes one or other of these roles to be more fundamental seems like a pretty esoteric question. But I've come to think that actually quite a bit turns on it. If you think beliefs are fundamentally the things that inquiry starts with, then there is a little gap in the argument that one should not inquire into what one already believes.

That argument, the one to the conclusion that one should not inquire into what one already believes, seems pretty simple. Assume one believes that p and is inquiring into the question p?. Our theory is that beliefs are appropriate starting points for inquiry, so it looks like this one should end pretty quickly. One can just argue p, therefore p, and close the inquiry. If the inquiry stays open longer than that, one is doing it wrong.

This looks like a pretty strong argument for a conclusion that a number of people have reached via different routes.²

If one knows the answer to some question at some time then one ought not to be investigating that question, or inquiring into it further ... at that time. (Friedman, 2017: 131)

There is something to be said for the claim that the person who knows they have turned the coffee pot off should not be going back to check. (Hawthorne and Stanley, 2008: 587)

Any such cases [of believing while inquiring] involve peculiarities (such as irrationality or fragmentation). (McGrath, 2021: 482n37)

So how could that argument fail? It could fail if there are reasons for adopting constraints on an inquiry. If there are reasons to not use all the tools at our disposal, there could be cases where an inquiry into p gets started, and we have reasons not to just say p, therefore p. At the highest possible level of abstraction, this doesn't sound very likely. It seems at

² These quotes were compiled by Elise Woodard (2020).

first like there should be something like a principle of total evidence for inquiry, saying that you can use whatever tools, whatever evidence, you have to hand. Such a principle, however, turns out to be false.

To warm up to this, consider an analogy to legal inquiries. There we are all familiar with the idea that some evidence might be inadmissible in some inquiries. Now the reasons for this are typically not epistemic. It's rather that we think the system as a whole will be more just if some kinds of evidence are excluded from some inquiries. That looks a bit different to the situation where an individual inquirer is just trying to find what's true. But we'll see that the analogy here is not quite as bad as it first looks.

In the rest of this section, I'll go over six kinds of cases where one can sensibly inquire into what one already knows. I don't think any of these examples constitute knock-down proofs of the possibility of rational inquiry into what one knows, and for reasons I'll get to later in the chapter, I don't really need them to. It is helpful to see the range of cases where inquiry into what one knows is useful.

5.1.1 Sensitivity Chasing

Guido Melchior (2019) argues that the point of *checking* is to establish a sensitive belief in the checked proposition. To motivate this, think about the following case. Florian has just weighed out the coffee beans for his morning pot of coffee. Naturally he uses the best scales he has for this purpose; it's important to get the coffee right. He starts wondering whether his scales have recently stopped being reliable. What does he do next? Here's one thing he doesn't do. He doesn't look at the beans on the scale, note that the scale says 24g, note that he knows they are 24g (via that excellent scale), and conclude that the scale is still working. That's no good at all; he has to use some other scale to check this one.

This is like the Problem of Easy Knowledge (Cohen, 2002), but note that it doesn't rely on the scale being a source of basic knowledge. Florian might have lots of independent evidence that the scale is good; it's from a good manufacturer and has been producing plausible results for a while. Still, if he wants to check it, he has to use something else. Here's the part that seems most surprising to me. Add to the story that he has a backup scale, one that he thinks is pretty good but not as good as his best scale. It's fine to use the backup scale to check the main scale, and not fine to use the scale to check itself. The best explanation for this is that checking requires sensitivity. Using the scale to test itself is a method that isn't sensitive to whether the scale is working. Using some other scale, even a less reliable one, to check whether it is working, is at least somewhat sensitive. Checking is, at least in part, a matter of *sensitivity chasing*. One reason it is often good to check what one knows is that sensitivity chasing is often sensible.

Sensitivity chasing is a perfectly acceptable goal in inquiry. One might inquire into p for the purpose of making one's belief in p more sensitive. Now assume, as most epistemologists believe, that one can know p even if one's belief is insensitive in various ways. One can know p even if one would still believe p were p false.³ If one has insensitive knowledge, it might be worthwhile to inquire into what one knows with the aim of generating sensitive knowledge. Indeed, this seems like a primary aim of what we call *checking*. Inquiring into p by saying p therefore p will not increase one's sensitivity to whether p is true. So it's worthwhile to not allow that move in the inquiry, if the aim is to increase sensitivity.

There are other examples that show the difference between knowing and checking. Slightly modifying an example from Frank Jackson (1987), imagine that someone wants to know what *The Age* said was the result of last night's game. One way to learn what *The Age* said would be to look up the result in *The Guardian*, and use one's background knowledge that they both report the same (correct) result. That's a way to come to know what *The Age* said. But it's not a way to check what *The Age* said. It's not a way to check because had *The Age* said anything different, you wouldn't have known. That's a kind of insensitivity. It's an insensitivity that's consistent with knowledge; one can know what a newspaper says by knowing the truth and that the newspaper reports the truth. This insensitivity is removed by proper checking. So checking aims for sensitivity that goes beyond belief, and beyond knowledge. Given that checking, i.e., chasing this kind of sensitivity, is rational, so is inquiring into what one knows.

³ One simple example from Saul Kripke (2011): I know that I do not falsely believe that I was born on the Galapagos Islands. But while this is knowledge, it is not a sensitive belief.

5.1.2 Rules

It's hard to always be perfectly rational. Sometimes it makes sense to not think too hard about things where getting the right answer would be quite literally more trouble than it's worth. I'll have much more to say about this point in Chapter 6, where I explore this insight from Frank Knight in greater depth:

It is evident that the rational thing to do is to be irrational, where deliberation and estimation cost more than they are worth. (Knight, 1921:67n1)

Knight is interested in the case where the rational thing to do is not inquire when inquiry would have minimal gains. There is another case that is more relevant here. Sometimes it is worth having a simple rule that says *Always inquire in these situations*, rather than having a metainquiry into whether inquiry is worthwhile right now. To make this a little less abstract, it might be worthwhile always checking that the door is locked when one closes it, even if one frequently knows that one has just locked the door. As Hawthorne and Amia Srinivasan (2013) point out, given the non-luminosity of evidence and knowledge, a simple rule like this might do better than any other realistic rule.

Often following rules about when to inquire will be part of one's professional responsibilities. I presented an example like this in chapter 7 of *Normative Externalism* (Weatherson, 2019): an inspector who is sent to do a random check of an establishment he had checked just a few days before. He knows everything is working well; he just checked it! But it's his job to check, and it's good to have random spot checks on top of regular checks, so it's good to run this inquiry. That's true even though the inspector knows how it will end.

5.1.3 Understanding

There is a famous puzzle about moral testimony. Something seems off about a person who simply believes moral principles on the basis of testimony, even from a trusted testifier. It's odd to convert to vegetarianism simply because someone you trust says that's what morality requires. There is also a famous answer to this puzzle, from Alison Hills (2009). (There are other answers too, including ones that deny the puzzle exists. To avoid going down too many rabbit holes, I'm going to assume for now the answer Hills gives is correct.) Hills says that moral testimony can give us moral knowledge, like any kind of testimony can provide knowledge, but it can't provide understanding. What's weird about the person who becomes a vegetarian on testimonial grounds alone is that they can't explain their actions, since they don't know why they are acting this way.

Beyond moral testimony, there seem to be many everyday cases of knowledge without understanding. One can know that Franz Ferdinand was assassinated in Sarajevo on June 28, 1914, without knowing why that happened. Or, indeed, one can know why one part of that is true, e.g., why it was that *Franz Ferdinand* was assassinated in Sarajevo on June 28, 1914, without knowing why he was assassinated *in Sarajevo*; or why he was assassinated *on June 28, 1914*. Given those facts, it is possible to seek understanding of something that one already knows.

In many cases, but not all, the search for understanding will look like a somewhat different inquiry to the search for knowledge. If one wants to know why Franz Ferdinand was assassinated *in Sarajevo*, one will inquire into the role that city plays in the history of relations between Austria-Hungary and Serbia. That will be a different kind of inquiry to determining whether the assassination really happened. But in the moral case things aren't this clear. Imagine again our person who hears from a trusted source that eating meat is wrong, but doesn't understand why this is so. They should do some moral inquiry. The inquiry will look, as far as I can see, very similar to the inquiry they will do in case they are working out whether eating meat is wrong. That is, it will look just like an inquiry into whether eating meat is wrong.

I think the best way to systematise things here is to take appearances at face value. Even once one is convinced eating meat is in fact wrong, if one doesn't know why it is, one will continue to inquire into the morality of eating meat. This inquiry is justified by the aim of coming to understand the wrongness of eating meat.

5.1.4 Defragmentation

Recall Professor Paresseux from Section 4.6.2. He's told that the visiting speaker this week is his old graduate school colleague Professor

Assidue. But he puts no effort into remembering this fact, and it slips from the front of his mind. The talk is approaching, and Paresseux wonders to himself, who's talking to us this afternoon? So he Googles the department talk schedule, sees that it is Assidue, and then says to himself "Ah, I knew that, I saw the email the other day."

It is very hard to fit the category of information that has 'slipped one's mind' into familiar epistemological categories.⁴ I think we should say that Paresseux is correct, and he did indeed know the answer to his inquiry before he started looking. After all, he could have retrieved the information by simply thinking hard about what had happened this week. The best explanation for why that's possible is that he did still know that Professor Assidue would be the speaker. But I also think it made sense for him to conduct an inquiry into this thing that he knew. It's much easier to Google something than to trawl one's memory for the answer, andmore reliable too. So this looks like a sensible inquiry for him to have conducted.

Following Andy Egan (2008), I treat this as a case where Paresseux's mind is 'fragmented', in the sense of David Lewis (1982) and Robert Stalnaker (1984). There is a part that contains the information about who the speaker is. That part isn't at the front of his attention, so he doesn't act on it. Still, it is a part of him; he knows that stuff. Still, it is better to conduct an inquiry, i.e., a Google search, than to rely on this knowledge. So it is rational to inquire into something one knows.

5.1.5 Public Reason

One unfortunate position an inquirer can find themselves in is knowing something is true, even understanding why it is true, and being unable to convince anyone of their result. At this point one needs more reasons, but where to find them? Often, the way to find them will be to do what anyone else would do if they were trying to find out if the thing itself were true. Here are two such examples, drawn from rather different parts of philosophy.

Michael Strevens (2020) argues that the effectiveness of science in the last 350 years is partially due to the fact that scientists have adopted

⁴ The point here is related to the discussion in Section 2.7.1 about how sometimes *knows* seems to just mean *possesses the information*.

an "iron rule": only empirical evidence counts. There are any number of ways one might come to rationally believe a scientific theory other than evidence. It might follow from broadly metaphysical principles one holds (at least in the early modern sense of metaphysical), it might be more elegant than any other theory, it might promise to unify seemingly disparate phenomena. But if you want to convince the scientific community, which includes both the collective body and the majority of individual scientists, you need data. So you go looking for data, even for theories you know are true on non-empirical grounds. Strevens thinks this is individually irrational, but collectively for the best. It's irrational for any one person to have just one way to come to believe things. But by incentivising the search for data in this way, we've collectively created an institution that has taken the measure of the world in ways previously unimaginable. There is something else valuable about data--it's available, at least in principle, to everyone. So even if you can't recreate my metaphysical intuitions, you can rerun my experiments. The iron rule doesn't just lead to more measurements being taken, it imposes a kind of public reason constraint on science. Only evidence that everyone can accept as evidence, and indeed that they could (at least in theory) create for themselves, counts.

This way of putting the point should remind us of an important strand in contemporary political philosophy, namely that political rules should satisfy a *public reason* constraint. As Jonathan Quong puts it

Public reason requires that the moral or political rules that regulate our common life be, in some sense, justifiable or acceptable to all those persons over whom the rules purport to have authority. (Quong, 2018)

Now as a matter of fact, we haven't had as much uptake of this meta-rule in politics as in science. But we can imagine a society where there is, in practice, a kind of public reason constraint. If you want your favourite rule to be part of the regulation of society, you have to come up with a justification of it that satisfies this constraint. In such a society, there will be people who have idiosyncratic ideas for rules that would be good rules for the community, ideas that they don't have public justifications for. In practice, the vast majority of these ideas will be bad ones. But some of them will not be. Indeed, a handful will even know that their ideas are good. Still, if this knowledge comes via idiosyncratic sources, they will need to come up with more public reasons if they want to see their rule implemented. As I suggested in the previous subsection, the way to find reasons for a moral claim is generally to inquire into whether that claim is true, or to at least act like that's what one is doing.

5.1.6 Evidence Gathering

In Section 9.6 I'm going to argue that having *p* as part of one's evidence might license inductive inferences that are not licensed by a smaller evidence set that doesn't include *p*, even if one knows *p* on the basis of that smaller set. If that's right, evidence gathering could be epistemically useful even if one already knows the evidence to be gathered.

5.1.7 Possible Responses

If this was a paper dedicated to proving that it is rational to inquire into what one knows, at this stage I'd have to show that a philosopher who denies that is ever rational has no good story to tell about these six cases. That would be a lot to show, since actually there is plenty that such a philosopher could say. They could deny that the inquiries are indeed rational. They could deny that the inquirers in question really do know the thing they are inquiring into, perhaps using IRT to back up that denial. They could deny that these are real inquiries, as opposed to some kind of ersatz inquiry. Or they could deny that this is really an inquiry into the very thing known, as opposed to an inquiry into some related proposition, like what the causal history of that thing was. They wouldn't even have to choose between these four; they could mix-andmatch to deal with the putative counterexamples.

At the end of the day, I don't think these responses will cover all the cases. But it would be a massive digression to defend that claim, and it isn't necessary for what's going to happen in the rest of this chapter. All I need is that there are people who very much look like they are conducting rational, genuine inquiries into things they already know. If there is a subtle way of explaining away that appearance, that won't matter for the story that's to come, since such subtleties will end up being good news for my side of the debate about IRT. The worry we're building up to is that IRT has no good explanation of what's happening in cases where someone seems to rationally, genuinely inquire into something they already know. If there are in fact no such cases, that can't be a problem!

One reason for thinking that some of these cases will work is that there is a fairly general recipe for constructing the cases. It's due to Elise Woodard (2020) and (independently) Arienne Falbo (2021). Start with the following two assumptions. First, inquiry is not just about collecting knowledge, but generally about improving one's epistemic position.⁵ Second, given fallibilism, one can know p but have a sub-optimal epistemic position. So one can know p, but (rationally) want to improve one's epistemic position with respect to p. If one acts to address that want, one will be inquiring into what one knows, and doing so rationally. Given IRT you should worry about whether every step in the last few sentences really does follow from the ones before it. But I suspect the general picture is right, especially, as Melchior (2019) stresses, in checks aimed at increasing sensitivity.

Looking ahead a little, the primary aim of the rest of the chapter will be to defuse some potential counterexamples to IRT that involve someone rationally inquiring, especially checking, what they know. My response will be disjunctive. Either inquiry solely aims at knowledge, or it does not. If inquiry does solely aim at knowledge, appearances in this case are deceiving, and the inquiry is not in fact rational. If, as I think, inquiry does not solely aim at knowledge, then the cases are not in fact counterexamples to IRT.

5.2 Using Knowledge in Inquiry

Sometimes an inquirer has reasons to deliberately hobble their own inquiry. They have reasons to conduct an inquiry with one hand tied behind their back. Perhaps those reasons come from the social norms of the enterprise they are engaged in, as Strevens suggests. Perhaps those reasons come from the fact that they are sensitivity chasing, as Melchior suggests, and only a restricted inquiry will increase sensitivity. Perhaps those reasons come from the fact that they are trying to follow rules, and the rules do not allow certain kinds of tools to be used. The unifying

⁵ When I say inquiry is about improving one's epistemic position, I don't mean that that's how inquirers represent what they are doing to themselves. That would be to over-intellectualise things. Rather, inquiry is about doing things that are, as a matter of fact, things that improve one's epistemic position. One can be improving one's epistemic position even if one self-represents one's actions in a more mundane way, e.g., as looking up when the coffee shop opens.

theme is that sometimes the inquirer wants not just to run an inquiry, but to run it in a particular way.

The core principle in my version of IRT is that someone who uses what they know in inquiry is immune to criticism on the grounds that what they are doing is epistemically risky. Equivalently, they are immune to criticism on the grounds that their premises might be false. That's compatible with saying that someone can know p, and be properly criticised for using p in inquiry. I motivated that restriction in Section 4.5 by looking at people whose use of p in inquiry can be criticised on relevance grounds. In this chapter we see several more reasons. Someone who has reasons to perform a restricted inquiry, especially someone whose aims can only be realised by conducting a properly restricted inquiry, can be criticised for overstepping those restrictions. That's fine, and totally consistent with IRT, as long as we pay attention not just to whether someone is being criticised, but why they are being criticised.

It isn't just my idiosyncratic version of IRT that escapes this criticism. Jeremy Fantl and Matthew McGrath defend a version of IRT that uses the following principle.

When you know a proposition p, no weaknesses in your epistemic position with respect to p—no weaknesses, that is, in your standing on any truth-relevant dimension with respect to p—stand in the way of pjustifying you in having further beliefs. (Fantl and McGrath, 2009: 64)

I'm going to come back in Section 9.9 to why I don't quite think that's right. But my disagreement turns on a fairly small technical point; I'm following Fantl and McGrath's lead much more than I'm diverging from them. These examples of properly restricted inquiry show how they too can accept rational inquiry into what one already knows.

Consider a person who is sensitivity chasing; they know p but want to have a more sensitive belief that p. So they conduct an inquiry into p, and reason to themselves p, therefore p. This closes the inquiry. Something has gone wrong. It isn't bad reasoning; one can't go wrong with identity. And it isn't that they use something they know as a premise; anything one knows can be used as a premise. It's that they had an aim that could only be met by a restricted inquiry, and they violated those restrictions. That's the incoherence here.

There is a way to read Fantl and McGrath's principle so that this case is a problem for them, but I don't think it's the right reading. The sensitivity of one's belief is, in their terms, part of the strength of one's epistemic position. So if one's belief was more sensitive, one wouldn't have a reason to be chasing sensitivity. So in this case, you might think it's weakness of epistemic position that's relevant; the weakness of epistemic position explains why the inquiry is being conducted in the first place. But I don't think that's fair. The principle only talks about how inquiry should be conducted, not about whether the inquiry should be conducted. So Fantl and McGrath could say, and I think this is the right way to interpret their position, that knowledge is compatible with the weakness in one's epistemic position being what explains why an inquiry is in order. It's just that knowledge is not compatible with that weakness preventing the application of the knowledge once the inquiry has begun.

5.3 Independence

These reflections on the nature of inquiry help tidy up a loose end from *Normative Externalism* (Weatherson, 2019). In that book I argued against David Christensen's Independence principle, but I didn't offer a fully satisfactory explanation for why the principle should seem plausible. Here's the principle in question.

Independence

In evaluating the epistemic credentials of another's expressed belief about P, in order to determine how (or whether) to modify my own belief about P, I should do so in a way that doesn't rely on the reasoning behind my initial belief about P. (Christensen, 2011: 1–2).

This is expressly stated as a principle about disagreement, but it is meant to apply to any kind of higher-order evidence. (This is made clear in "Formulating Independence" (Christensen, 2019), which also includes some new thoughts about how Christensen now thinks the principle should be stated.) I argued that this couldn't be right in general; it gives the wrong results in clear cases, and leads to regresses. Still, it seems plausible that something like this should be right. In *Normative Externalism* I hinted at an inquiry-theoretic proposal about what that similar truth might be. (See, for example, the response to Clayton Littlejohn (2018), at the top of page 178.) But I never really spelled it out. Here's what I now think the right thing to say is.⁶

Peer disagreement, or really any other kind of higher-order evidence, gives a thinker a reason to conduct an inquiry into whether their earlier thinking was correct. Further, it gives them reason to conduct an inquiry that is restricted in a particular way. The restriction is that they should not rely on the reasoning from their earlier thinking. Putting those two things together, we get that disagreement about p gives someone who believes p reason to inquire into p using a different approach, any different approach, from what they previously used.

Once we've got a principle about reasons, we could try formulating this as a defeasible rule. It's plausible that one should adopt the defeasible rule of conducting such an inquiry whenever one sees a disagreement, or some other kind of potentially defeating higher-order evidence. As long as one builds enough into the defeasibility clause, such a rule won't be subject to the counterexamples I described, or the ones that have caused Christensen (2019) to have second thoughts about the right formulation of the rule. After all, every counterexample will naturally fall into the defeasibility clause.

Such a rule could be justified by the observation that it will probably be beneficial in the long run for people like us to adopt it. Double checking isn't that hard, and can be very useful. Getting stuck in a bad epistemic picture can have devastating consequences; it's good to step back from time to time to see if that's happening to us. Disagreements with peers are a natural trigger for that kind of inquiry. Those same benefits can explain why disagreement, or other kinds of higher-order evidence, give us reason to double check.

But why should one conduct a restricted inquiry here? Given the stakes, we're trying to work out whether we've got ourselves into a bad epistemic state, shouldn't we throw everything we have at the problem? That would be bad, since Independence expressly bars the thinker from using some of the tools at their disposal. It requires them to not do the same kind of inquiry they did before, which presumably was the

⁶ The picture I'm about to give is really similar to the one laid out by Andy Egan (2008). We're interested in different kinds of cases, but the idea that a cognitive system might work best by allowing one part to check on another using just the evidence the first part has endorsed is one I'm just taking from him. If I'd seen this connection when writing *Normative Externalism* I would have connected it to the discussion of Madisonian moral psychology in part I of that book.

one they thought best suited to the problem. That's a big restriction, and needs some justification. I can offer two kinds of justification, not entirely distinct.

The point of having a rule like this, a rule like *Double check your reasoning when a peer disagrees,* is to prevent us falling into epistemic states that are local but not global equilibria. The states we're worried about are ones where any small change will make the epistemic state worse, but large changes will make things better. Picturesquely, we've reached the top of a small hill when we want to climb a mountain. We should be somewhere higher, but any step will be downhill. It's good to not get stuck in places like this, and nudges from friends are a way out.

If we want to check whether we're in such a bad situation, we want a test that is sensitive to whether we are. That is, we want a test that would say something different if we were in that situation to what it would say if we were doing well. (This is Melchior's point about the aim of tests.) Just conducting the same inquiry we previously conducted will typically not be sensitive in this way. Or, more precisely, it will be sensitive to something like performance errors, but not competence errors. We need something more sensitive if the aim is to avoid getting stuck in local equilibria, and that requires setting aside the work we've previously done.

One of the reasons that local equilibria can be sticky is that we know our way around them well. We know all the ways in which one part of the picture we have supports the other parts. We typically don't know how to think about other pictures so clearly. We don't know, don't see, the ways in which other pictures might 'hang together' as well as ours does. We are inevitably going to be biased towards our own ways of thinking. So it's worthwhile to try to level the playing field, by looking at how things would seem if we didn't have our own distinctive way of thinking.

None of this is to take back anything I said in *Normative Externalism*. Disagreement with a peer known to have the same evidence does not give someone a reason to reject a well-formed belief. It gives them a reason to double check that belief. As I've been stressing all chapter, one can double check one's beliefs, and even one's knowledge. That is what should happen here.

Finally, thinking of disagreement as providing a reason to double check provides a nice explanation of one of the harder examples in *Normative Externalism*, the case of Efrosyni (see Weatherson, 2019: 222).

She does a calculation, then double checks it by a different technique, then hears that a peer disagrees. What should she do now? I think typically she should do nothing. The disagreement gives her a reason to double check each calculation she did, but she's already carried out that double check. This is, I think, the intuitively right result. If someone has already double checked their work, they typically do not need to check again. Perhaps in some rare case they could get reason to double check the 'combined' inquiry, consisting of the initial inquiry plus the double check. But that's rare; usually they should respond to the disagreement by showing their work.

With this picture of the relationship between knowledge, inquiry, and checking in place, it's time (at last) to return to potential counterexamples to IRT.

5.4 Double Checking

In her 2008 paper "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning", Jessica Brown (2008) runs through a bunch of cases where, she says, intuitively someone knows a proposition but they cannot use it in practical deliberation. The first of these cases has been frequently cited as a problem for the kind of view I'm defending.

A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is lying anaesthetised on the operating table. The operation hasn't started as the surgeon is consulting the patient's notes. The student is puzzled and asks one of the nurses what's going on:

Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn't operate before checking the patient's records.

I think there are pretty good arguments that checking the chart is the right thing to do even if the surgeon knows which kidney is diseased, so this case isn't a problem for the views about knowledge and action that I'm defending.

In medical contexts, intuitions about appropriate action very rarely track expected utility maximisation.⁷ This is one reason why it is so easy to come up with medical counterexamples to act utilitarianism for introductory ethics classes. Instead, intuitions about appropriate actions here are more likely to align with rule utilitarianism. The rule *Double check the notes before removing an organ* seems like it will on average maximise utility, even if it would not help in this case.

To connect this to the discussion in Section 5.1, the surgeon here is doing a bit of mostly harmless sensitivity chasing. Before checking the notes, their belief that the left kidney was diseased was not sensitive to the possibility that they'd misremembered the morning meeting; after checking the notes it is. Since busy surgeons do sometimes misremember meetings some hours earlier, this is a reasonable bit of sensitivity for the surgeon to chase, and for the rule-makers to require be chased.

All that can be true even if the surgeon knows which kidney is diseased. If they inquired into which kidney should be removed, and used their knowledge about which kidney was diseased, they would get the right answer. In some sense this would be a perfectly conducted inquiry. But it would not be an inquiry that delivered what the surgeon was looking for, and what the regulators require them to look for: a belief that was sensitive to the possibility of an error in memory.

These considerations don't just defend IRT against the example, they show how IRT can be used to resolve a puzzle about a related case. Continue Brown's story by imagining that every time the surgeon raises the scalpel to make the first incision, they instead go back to look at the notes to check they are removing the correct kidney. Now we have the following conversation.

Student: I don't understand. Why is she looking at the patient's records for the seventeenth time? She just looked at the notes each minute for the last sixteen minutes; she knows which kidney it is. *Nurse*: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn't operate before checking the patient's records. (Brown, 2008: 176)

This is a really bad defence of the surgeon's actions. We are owed a story about why it is a bad defence. My story starts with the point that Student

⁷ Jonathan Ichikawa (2017: 152ff) makes this point well in responding to Brown.

is right to ask why she is inquiring into something she knows. While as we've seen there are cases where that is appropriate, these cases are somewhat unusual. It's a reasonable default assumption that inquiry into something one knows is mistaken. That assumption is only defeated if there is some other worthwhile epistemic good that can be attained. In this case, there isn't, since sensitivity to whether one misread the chart the last sixteen times isn't a worthwhile kind of sensitivity to get.

In general, anyone who wants to separate out knowledge from action, and do so on account of the fact that sometimes we double check things we know, owes a story about why we don't also triple-check, quadruplecheck, and so on. I suspect such a story won't be easy to tell.

Brown has another example that hasn't attracted nearly as much attention in the literature. This is unfortunate since I think it's a more pressing problem for the view Brown is attacking.

A husband is berating his friend for not telling him that his wife has been having an affair even though the friend has known of the affair for weeks.

Husband: Why didn't you say she was having an affair? You've known for weeks.

Friend: Ok, I admit I knew, but it wouldn't have been right for me to say anything before I was absolutely sure. I knew the damage it would cause to your marriage. (Brown, 2008: 176–177)

In this case, the tricks I was deploying in Section 5.1 don't seem to help. There is no further epistemic good that Friend obtains by waiting further.

That said, my intuition here is that Friend's speech is just incoherent. Or, at least, it is incoherent if we take the final statement at face value. I think we shouldn't do that; Friend didn't really know about the affair.⁸

There are two things that might be going on in this case. My best guess is that the explanation for why Friend's statement seems so natural relies on both of them.

⁸ The particular versions of IRT Brown was responding to in the 2008 paper were heavily motivated by intuitions about cases. Brown argues, quite correctly I think, that those theories aren't entitled to appeal to arguments that the intuitions which go against them are mistaken. After all, if IRT is just motivated by intuitions, the argument that knowledge is not sensitive to interests is just as good an argument against those intuitions as the arguments that IRT defenders can make about this example. Happily, my version of IRT is not motivated just by intuitions about cases, so I don't have to worry about this dialectical point.

First, we do sometimes use 'know' in a purely informational sense. We saw this in Section 5.1.4 with Paresseux's claim that he *knew* Assidue was visiting. He possessed the information, though little more than that. Still, in context this can be enough to ascribe knowledge.

Second, we can be very flexible about past-tense knowledge claims when we, the current speakers, know how things turned out. After our sports team loses a game they should have won, we might say "I had a bad feeling about today, I knew we were going to mess it up". In most cases it would be weird to say the speaker even thought their team would mess up, let alone believed it. (Why didn't they bet on the opposition if they thought the result was a foregone conclusion?) But even if they did believe it, we really don't think *bad feelings* are appropriate grounds for knowledge. And yet, the speakers claim that they knew the team would mess up sounds fine.

Is Friend's statement like Paresseux's knowing (i.e., possessing the information) that Assidue would be visiting, or the sports fan's knowing (i.e., having an accurate premonition) that their team would mess up? My guess is that it's a bit of both. Either way, the Friend didn't know, in the sense of *know* relevant to epistemology, about the affair.

The general methodological point is that these last two senses of knowledge do seem different to what we typically talk about in epistemology. It's possible, as I noted in Section 2.7.1, that considering the information-possession sense of knowledge is important for thinking through whether any kind of contextualism is true. I don't think the 'bad feeling' cases are relevant to anything in epistemology, save for cases where we might need to explain away intuitions that they are involved in. Maybe that's what's happening in Brown's second case.

5.5 The Need to Inquire

So far I've mostly talked about inquiries that a person is actually conducting. But we should also think about the inquiries that they should conduct. Consider the following two abstractly described possibilities.

A person believes p for good reasons, and it is true, and there are no weird things happening that characterise typical gaps between rational true belief and knowledge. There is some action φ they are considering that will have mildly good consequences if p, and absolutely catastrophic consequences if $\neg p$. One of the alternatives to φ is first checking whether p, which would be trivial, and then doing φ iff p. We've seen lots of these cases before, but here's the new twist. The person absolutely does not care about the catastrophic consequences. They will all fall on people the person could not care less about. So they are planning to simply do φ , for the good consequences. Since p is true, nothing bad will happen. Still, it seems something has gone wrong. We want to say that they've been reckless, that they've taken an immoral risk. But it isn't risky to do something that you know won't have bad consequences. So they do not know that p, and for similar reasons to why Anisa doesn't know that p. Yet the version of IRT that I've given so far doesn't say that they don't know that p.

The second case has the same initial structure as the first. The person believes *p* for good reasons, it's true, and there is no funny business going on-no fake barns or the like blocking knowledge. They are thinking about doing φ . They know that if *p* is true, φ will have a small benefit. They also know that it would be completely trivial to verify whether *p* is true. They also in some sense know that if they do φ , and p is false, it will be absolutely catastrophic. And they care about the catastrophe. But they've sort of forgotten this fact about φ . It's not that it has totally vanished from their mind. But they aren't attending to it, and it doesn't form any part of their deliberation when thinking about φ . So they do φ , nothing bad happens, and later when someone asks them whether they were worried about the possible catastrophe, they are shocked that they would do something so reckless. They are shocked, that is, that they forgot that it was important to confirm whether *p* was true before doing φ . It feels, from the inside, like they got away with taking a terrible risk. But if they knew *p*, it should not seem like a risk, it should seem like rational action. (Just like they would think doing φ after checking whether p was rational action.) So this too should be a case where we say knowledge fails for practical reasons. (I'm going to come back to a version of this case in Section 8.1, where it will be useful for highlighting one of the few points where I disagree with the theory that Jeremy Fantl and Matthew McGrath (2002, 2009) endorse.)

The natural thing to say here is that in each case, the person should conduct an inquiry. They should check whether p is true. In that inquiry, they shouldn't take p for granted. They shouldn't take it for granted for

a very particular reason, because it might be false. If they knew *p*, they could take it for granted, or, at least, if they couldn't, it would be for some reason other than that *p* might be false. So they don't know that *p*.

What these two types of case show is that knowledge is not just sensitive to what one is actually inquiring into, it is also sensitive to what one should be inquiring into. If one should inquire into Q, and were one to inquire into Q, one shouldn't take p for granted because it might be false—one doesn't know p.

This is a kind of moral encroachment in the sense of Rima Basu and Mark Schroeder (2019). What one knows might be sensitive to one's moral obligations in inquiry. Imagine two people both take p for granted in making a decision that affects other people. This is mostly fine because p is true, and they had good reasons to take it for granted. Still, there was some risk to others, and they could have checked whether p was actually true before acting, but in each case they had other things they would rather be doing than checking p. What differs between the two people is what they would rather be doing. The first could have checked, but it would have taken them away from a rescue operation in progress; the second could have checked, but it would have taken them away from their social media feed. If the theory I've developed so far is correct, then the first knows that p, and the second does not, and the difference comes down to the differing moral importance of contributing to rescue operations and social media.

It's worth recalling here that the methodology I'm using in this book is perhaps a little different to a common methodology in this area. I don't think that if you fill out the two cases from the last paragraph in full detail, it will be intuitively obvious that one person knows and the other doesn't, and that's evidence for IRT. Rather, I think that it's plausible that one isn't being reckless by acting on what one knows, and this principle, combined with anti-sceptical principles and judgments about which acts are indeed reckless, leads to IRT. As always, these cases allow for four broad classes of response: the sceptic who denies there is knowledge even in the low-stakes case; the epistemicist who denies the intuitions about which actions are reckless; the orthodox theorist who says that acting on what one knows can be reckless; and the pragmatist, who accepts both the intuitions about which acts are reckless and how knowledge connects to recklessness, and infers that knowledge is sensitive to pragmatic, and in this case moral, factors.

5.6 Multiple Inquiries

IRT says that what one knows depends on what one is inquiring into. It would be very convenient if there was a position in the logical form of knowledge ascriptions for inquiries. That is, it would be very convenient if the logical form of *S* knows that *p* was something like *Ktspi*, where *t* is the time, *s* is the knower, *p* is what's known, and *i* is the inquiry it is known in. Then we could say that one condition on such a knowledge claim being true is that at *t*, *s* can properly use *p* as a starting point in inquiry *i*.⁹ Unfortunately for IRT, that's not the logical form of knowledge ascriptions. The *t*, *s*, and *p* are there all right, but not the *i*. Fortunately for IRT, the logical form does have reference to a knower, that *s*. Since knowers undertake inquiries, we can bring in the inquiries via the knower. All knowledge is inquiry-relative, we say, and it is relative to the inquiries of the person to whom knowledge is being ascribed.

If every person was, at each time, undertaking precisely one inquiry, everything would fall into place very nicely. Given *t* and *s*, we could guarantee the unique existence of an *i*, and it would be as if there was an *i* in the logical form, as IRT would like. Unfortunately, that's not close to being true. Some people at some times are making no inquiries, e.g., when they are asleep. And some people at some times are making many inquiries. The former case is no problem for IRT. If the person is making no inquiries, then what they know is determined by 'traditional' factors, such as what they believe, whether those beliefs are true, grounded in the evidence, safe, and so on. The case where someone is engaged in multiple inquiries is a little harder.

The view I'll defend is that the person knows p only if p can properly be used as a starting point in all the inquiries the person is engaged in. This has a surprising, and not entirely welcome, side effect. It means that some people don't know p, and hence can't use p in an inquiry i, even though they could use p as a starting point to i if i were the only inquiry

⁹ More precisely, as I said in Section 4.5, if they use *p* in *i*, that won't be subject to criticism on the grounds that *p* might be false. I'll use the more informal version in the text in what follows to increase readability.

they were engaged in. This is a somewhat more sceptical result than I like, but I suspect it's the best choice out of a bad lot. The only other options I can see are to either try to find ways to get *i* back into the logical form of knowledge ascriptions, or to adopt a novel form of relativism that says knowledge claims are true or false relative to inquiries, or to say that the person conducting multiple inquiries is fragmented, and each of the fragments has their own knowledge. None of these moves strikes me as remotely plausible, and so we're forced to have some kind of view where we quantify over the inquiries a person is engaged in.

In the rest of this section, I have three aims. First, to make what I've said so far less abstract, by describing a case where someone has multiple inquiries, and this matters in surprising ways. Second, to say why it isn't great that IRT is forced to say that someone doesn't know something that is otherwise usable in an inquiry they are engaged in. Third, to say why this isn't a devastating result, even though it's not exactly a happy one.

Our example of someone with multiple inquiries will be a historian called Tori. She has been taught, like everyone else, that the Battle of Hastings was in 1066. For most purposes she takes that to be one of the fixed points in the historical record. But she's noticed some anomalies in some of the documents from around that time, anomalies that would be explained by the battle being in 1067. She's seen enough documents to know that the overwhelming likelihood is that these anomalies have some simple explanation, like a transcription error. But in her spare time over the last few years, she has been investigating off and on whether the best explanation might be that everyone else has the date of the battle wrong, and in fact it was in 1067.

If it is worth inquiring into the date of the Battle of Hastings, it is not sensible to take the date of the battle as fixed. That would make the inquiry very short. So if it's reasonable for Tori to conduct this inquiry, then while she is conducting it, she does not know when the Battle of Hastings took place.

If this inquiry into the date is something she has been working on in her spare time for years, she has presumably had other jobs that did not involve trying to overturn the historical record about one of the central events in British history. In some of those jobs, it will have been sensible to take as given when the Battle of Hastings, and hence the beginning of Norman rule over Britain, took place. So there will be contexts, ones where her primary focus is on an everyday question where one takes for granted the common assumptions about British history, but she still has as a background project this idea that maybe the Battle of Hastings took place a year later, where IRT seems to get into trouble. It wants to say that for the purposes of her everyday inquiries, Tori knows the Battle of Hastings took place in 1066. After all, this is a true, rational belief, that is based in the right way in the facts, and which is reasonably taken as a starting point for this very inquiry. That looks like, relative to that inquiry, it is knowledge. But for the purposes of finding out the best explanation of the anomalies, she does not know when the battle took place, on pain of not being able to rationally investigate one possible explanation.

My version of IRT says that knowledge is relative to inquirers, not to inquiries, so I can't say that she knows the date relative to one inquiry but not another. That's not great. In the everyday inquiry Tori is exactly like someone who knows when the Battle of Hastings was in what look like all the relevant features, and yet she doesn't know. How can we explain away this anomaly?

The first thing to note is that even if Tori loses the knowledge that the Battle of Hastings was in 1066, she keeps her voluminous evidence that the Battle happened then. In most inquiries, anything she might infer from a claim about the Battle's date, she can infer from that evidence. So she'll still, on the whole, be able to draw the same conclusions in other inquiries as if she kept that knowledge.

Usually there are two reasons for keeping the conclusions of one's inquiries and not one's evidence. First, it helps with clutter avoidance (Harman, 1986: 12). If a knowledge of history required knowing not just a bunch of things about what happened, when it happened, and ideally why it happened, but also knowing how and where one learned these facts, then even the most basic knowledge of history would be beyond most of us. Second, it makes certain kind of inferences much smoother to go through various steps rather than applying something like cut-elimination and getting rid of the middle steps. That is, it's easier for Tori to infer from some evidence that the Battle of Hastings was in 1066, and then from that and some other evidence to draw further conclusions, than it is to draw inferences directly from the underlying

evidence. But while both of these considerations are very powerful ones in general, one would definitely not like to never store or rely on intermediate conclusions in inquiry, they aren't nearly as powerful in any specific case. If there's one step in an inquiry that one is unsure of on other grounds, it's not a huge effort to retain one's evidence for *that step*, and replace inferences that rely on it with inferences that rely on the underlying evidence.

The other thing to note is that we can explain Tori's behaviour in inquiries without positing more knowledge to her than IRT allows. The key thing is to replace the familiar Knowledge Norm of Assertion with the slightly more complicated Sufficient Evidence Norm of Assertion.

Knowledge Norm of Assertion

One must: Assert *p* only if one knows that *p*.

Sufficient Evidence Norm of Assertion

One must: Assert *p* only if one's evidence is sufficient for one's audience to know that *p*.

If one identifies evidence with knowledge, then it's hard to see any space between these two. I don't quite endorse that identification for reasons that I'll go over more in Chapter 9, but I mention it here just to note that this need not be a radical revision.

If the norms do come apart, then the latter seems to play better with IRT. Imagine that *S* is talking to some people who are facing a long-shot bet on whether *p*. These people would not be best off, in expectation, taking *p* for granted. Unfortunately, *S* doesn't care about the welfare of these people, though for some reason they do care about being a good informant and testifier. Further imagine that *S*'s evidence for *p*, while strong, isn't quite strong enough to justify the audience in taking this long-shot bet. Then it is wrong for *S* to simply say that *p*.

The picture behind the Sufficient Evidence Norm of Assertion is that one should say p only if one's audience can take p as a starting point in inquiry. Sometimes one might violate this norm without being subject to blame, as when it turns out one's audience has an unexpectedly long-odds bet on p. In normal cases, however, where one knows at least something about one's audience, one should calibrate one's assertions to the projects of one's audience. This picture seems to get two possible cases where Tori is involved in a group inquiry just right.

In the first (more normal) case, Tori is working with a group of people who do not share her worry about the anomalies in the dating of the Battle of Hastings. They think the date is a settled fact. In their presence Tori can speak as if it is settled. After all, her evidence suffices for her audience to know when the Battle was, given their lack of interest in odd anomalies.

In the second (somewhat odder) case, Tori is working with a group of people, one of whom shares her concerns about these anomalies. In the context of the other inquiry (i.e., not the inquiry into the date of the Battle), Tori says, "The Battle of Hastings was in 1066." It would be reasonable for the other person who shares her concerns about the anomalies to conclude that Tori had satisfied herself that the anomalies were just mistakes, and the Battle really was in 1066. That's because, I say, the unqualified assertion would be improper unless Tori had resolved these concerns to a standard that would be satisfying to the two of them. This case is a bit odd, as it does require the coincidental presence of two people with unusual interests, but I think the Sufficient Evidence Norm plus IRT gets them right.

Two final notes about this case.

First, I've crafted the Sufficient Evidence Norm to be the variation on the Knowledge Norm that a defender of IRT should like. But one might suspect the Knowledge Norm on independent grounds, e.g., because it gets the cases in Maitra and Weatherson (2010) wrong. I think the Sufficient Evidence Norm should be tinkered with to handle those cases, but I'm not exactly sure how this should go. Still, the tinkering shouldn't undermine the way IRT handles these cases.

Second, there is a really interesting historical question around here. Imagine you have a community that governs itself by the Sufficient Evidence Norm. And then someone comes along and invents the scientific journal, and all of a sudden it's possible to assert things with no knowledge of what is at stake for one's audience. How should one react, especially given the usefulness of the scientific journal for conducting inquiries that are widely distributed over space and time? A natural move would be to develop some new interest-invariant standards for printed assertion, and hopefully make it clear to both writers and readers what these standards are.

Once upon a time I had hoped this book would include an argument that the development of interest-invariant epistemology was just such a reaction to the invention of the printing press and, somewhat later, to the adoption of scientific journals as important conduits for sharing information in distributed inquiries. I still think something like this is arguably true, at least if we mean the development of interest-invariant norms for what I called in Chapter 1 'sub-optimal' epistemology. But defending this claim would require a different book, and a writer with very different skills, to this one.

So I'll just leave this as a conjecture for future research. What most philosophers call 'traditional' epistemological views, i.e., fallibilism plus interest-invariance, might just be a response to a relatively recent technological innovation.