

The background of the entire image is a photograph of a sunset over the Golden Gate Bridge in San Francisco. The sky is a mix of deep purple, blue, and orange, with wispy clouds catching the low light. The bridge's silhouette is visible against the horizon, with its iconic towers and suspension cables. The water in the foreground reflects the warm colors of the sunset.

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All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at 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-suspension-bridgeduring-golden-hour-Vnz7o7LeuDs>.

Cover design: Jeevanjot Kaur Nagpal

7. Changes

My version of IRT version shares defects with more familiar versions of IRT. For instance, it is subject to the criticism that Crispin Wright makes here.

[A] situation may arise ... when we can truly affirm an 'ugly conjunction' like:

X didn't (have enough evidence to) know P at t but does at t*
and has exactly the same body of P-relevant evidence at t* as at t.

Such a remark seems drastically foreign to the concept of knowledge we actually have. It seems absurd to suppose that a thinker can acquire knowledge without further investigation simply because his practical interests happen so to change as to reduce the importance of the matter at hand. Another potential kind of ugly conjunction is the synchronic case for different subjects:

X knows that P but Y does not, and X and Y have exactly the same body of P-relevant evidence.

when affirmed purely because X and Y have sufficiently different practical interests. IRI, as we noted earlier, must seemingly allow that instances of such a conjunction can be true. (Wright, 2018: 368)

That's right; I do allow that instances of such a conjunction can be true. A similar objection has been made by Gillian Russell and John Doris (2009), by Michael Blome-Tillmann (2009), and by Daniel Eaton and Timothy Pickavance (2015). My main reply to these objections is that they over-generate and would be successful objections to any theory that separates knowledge from rational true belief. Since knowledge does not equal rational true belief, no such objection can work.¹

¹ This reply was first made in my (2016b), and earlier replies to Russell and Doris, as well as Blome-Tillman, were made in my (2011), although I now believe that

7.1 Overview of Replies

I'm going to quickly go over five responses to this objection. I think at some level all five are correct. The first two, however, would probably do little to persuade anyone not already committed to IRT. The last three are more persuasive, and I'll develop each of them in a subsequent section.

The first thing one could say about these objections is that they simply highlight a prominent feature of the view—namely, that it allows knowledge to turn on non-alethic features—and then object to that very feature. As a result, the objections are blatantly question-begging. It doesn't accomplish much more than pointing out the obvious. The opponents think that this view is radical. And of course the objections to radical views will end up being question-begging (Lewis, 1982). Saying that one's opponents are begging the question might make you feel better—you don't have to be persuaded by their arguments—but doesn't actually move the debate forward. We can, and must, do better.

A second thing to say is that on some versions of IRT, it will be very hard to state the objection. Consider a version of IRT that also accepts $E=K$, the thesis that one's evidence is all and only what one knows. This is hardly an obscure version of the view; it's what is defended by Jason Stanley (2005). Now it will not be true, according to such a view, that there are, as Wright suggests, two people who have the same evidence but different knowledge. That's impossible, since having different knowledge literally entails, in this view, that they have different evidence. But does this make the objection go away, or does it just make it harder to state? I'm mostly inclined to think it's the latter. There is still something weird about people who have the same input from the world, and the same reactions to that input, but who differ in what they know about the world. So this response, while more useful than the last one, i.e., not totally useless, won't quite work either.

A third response challenges head on the intuition about 'weirdness' mentioned in the previous paragraph. One of the consequences of the vast Gettier literature is that there are any number of cases where people have the same inputs, the same true beliefs based on those inputs, but different knowledge. It's trivial to get these inter-world versions of a case

those replies did not quite get to the heart of the matter.

like this, and maybe that's enough to undermine the intuition. More generally, it's hard to state, and endorse, the intuition that the interest-relative theory is flawed without implicitly committing oneself to something closely resembling the JTB theory of knowledge. And since that theory is false, that's kind of bad news for that intuition. Or, perhaps more carefully, either that theory is false, or justification is understood in terms of knowledge, as on the $E=K$ picture. And appealing to $E=K$ might be an independent way to respond to the challenge. I'll spell out this response more fully in Section 7.2.

A fourth response aims to undermine the intuition in a different way. There is something fundamentally right about the JTB theory of knowledge, at least if we don't presuppose that the justification, the *J*, gets an internalist spin. But it can't be that the theory is extensionally correct. What is it? My conjecture is that knowledge is *built*, in the sense described by Karen Bennett (2017), out of those three components—justification, truth, and belief. Now this needs a notion of building that doesn't involve necessitation, and spelling that out would be a task for a different (and longer!) book. I'll try and say enough in Section 7.3 to make it at least minimally plausible that this conjecture is true, and that it is consistent with IRT.

The fifth response, and the one I want to lean on the most, comes from Nilanjan Das (2016). On the most plausible ways of articulating what the differences are between JTB and knowledge, it's not just that the differences will depend on 'non-standard' factors, it's that they will often depend on interests. Whether a belief is safe, or sensitive, or produced by a reliable method, or apt, or virtuous, or any other plausible criteria you might want, depends in part on the interests of the believer. More carefully, whether a belief satisfies any one of those properties can be counterfactually dependent on the interests of the believer. So I conclude these objections massively over-generate. If they are right, they show that practically every theory of knowledge produced in the last several decades is false. But it's really implausible that these kinds of considerations could show that. So the objection fails. I'll end in Section 7.4 by spelling out this response.

7.2 So Long JTB

The story of investigations into knowledge over the last sixty years is the story of making the list of things knowledge is sensitive to ever longer. The thesis of this book is that human interests, in particular the interests of the would-be knower, should be added to that list. But to defend that thesis, and especially to defend it from the kind of blank stare objection that I'm worrying about in this chapter, it helps to have the list in front of us. So I'm going to describe a mundane case of knowledge, then discuss various ways in which that knowledge could be lost if the world were different.

Our protagonist, Charlotte, is reading a book about the build up to World War One. In the base case, the book is Christopher Clark's *The Sleepwalkers* (Clark, 2012), though in some of the variants we'll discuss she reads a less impressive book. In it she reads the remarkable story of Henriette Caillaux, the second wife of anti-war French politician Joseph Caillaux. As you may already know, Henriette Caillaux shot and killed Gaston Calmette, the editor of *Le Figaro*, after *Le Figaro* published a string of damaging articles about Joseph Caillaux. The killing took place on March 16, 1914, and the trial was that July. It ended on July 28 with her acquittal.

Charlotte reads all of this and believes it. And indeed it is true. And the book is reliable. Although Charlotte does believe what the book says about Henriette Caillaux, she is not credulous. She is an attentive enough, and skilled enough, reader of contemporary history to know when historians are likely to be going out on a limb, and when they are not being as clear as one might like in reflecting how equivocal the evidence is. But Clark is a good historian, and Charlotte is a good reader, and the beliefs she takes from the book are both true and supported by the underlying evidence.

Focus for now on this proposition.

Henriette Caillaux's trial for the murder of Gaston Calmette ended in her acquittal in late July 1914.

Call this proposition *p*. In this base case, Charlotte knows that *p*. But there are ever so many ways in which Charlotte could fail to have known it. The following three are particularly important.

Variant J

Charlotte didn't finish the book. She only got as far as the start of Caillaux's trial, but lost interest in the machinations of the diplomats in the late stages of the July crisis. Still, she had a strong hunch that Caillaux would be acquitted and, on just this basis, firmly believed that she would be.

Variant T

Charlotte is in a world where things went just as in the actual world up to the trial, but then Caillaux was found guilty. Despite this, Charlotte reads a book that is word-for-word identical to Clark's book. That is, it falsely says that Caillaux was acquitted, before quickly moving back to talking about the war. Charlotte believes, falsely, that *p*.

Variant B

Charlotte reads the book to the end, but she can't believe that Caillaux would be acquitted. The evidence was conclusive, she thought. She is torn because she also can't really believe a historian would get such a clear fact wrong. But she also can't believe anyone would be acquitted in such a trial. So she withholds judgment on the matter, not sure what actually happened in Caillaux's trial.

Charlotte does not know that *p* in all three scenarios. These cases are good evidence that knowledge requires justification, truth, and belief. In Variant J, Charlotte's belief in *p* is not justified, but rather a mere hunch, so she doesn't know. In Variant T, Charlotte's belief is incorrect, making it an honest mistake and hence not knowledge. In Variant B, Charlotte lacks knowledge because she doesn't even believe *p*; she has the evidence, but does not accept it.

There are philosophers who argue that the conditions in all three cases are not strictly necessary. However, I won't be discussing these points as it would take us too far afield. Instead, I'll assume that Variant J demonstrates the need for justification or some form of rationality for knowledge. Variant T shows that knowledge requires truth, and Variant B shows that belief or strong acceptance is necessary for knowledge.

For a short while in the mid-20th century, some philosophers thought these conditions were not merely necessary for knowledge, but jointly sufficient. To know that *p* just is to have a justified, true belief that *p*. This became known, largely in retrospect, as the JTB theory of knowledge. It fell out of fashion dramatically after a short but decisive criticism was published by Edmund Gettier (1963). But Gettier's criticism was

not original; he had independently rediscovered a point made by the 8th-century philosopher Dharmottara (Nagel, 2014). Here is a version of the kind of case Dharmottara discovered.

Variant D

Charlotte stops reading before the denouement. She thinks Caillaux was acquitted, not on a hunch, but because she read in another book that official France was too disorganized in July 1914 to convict any murderers. This is untrue, but Charlotte used it to arrive at the correct conclusion that *p*.

In Variant D, Charlotte lacks knowledge of *p* because basing one's reasoning on a falsehood typically does not establish knowledge. So whether one knows is influenced by the accuracy of the grounds for one's belief. The subsequent variations may not be as straightforward, as determining whether Charlotte knows *p* will be more controversial. They are all instances where it is plausible that knowledge is sensitive to more factors than we've seen so far. The first case is a version of an example from Gilbert Harman (1973: 143ff).

Variant H

Charlotte's unfamiliarity with Henriette Caillaux is surprising, because in her world Caillaux is as infamous as killers like Ned Kelly, Jack the Ripper, and Lee Harvey Oswald. Her killing of Calmette has been the subject of numerous novels, plays, and movies. But all these renditions have a fictionalised ending: Caillaux is convicted and executed. The authorities were so embarrassed by the actual ending of the trial, where Caillaux was acquitted, that they successfully conspired to convince the public that this never happened. Charlotte, coincidentally, is the only person who hasn't heard of Caillaux's story. When she reads a word-for-word copy of Clark's book, she doesn't realize it's controversial and believes that *p*. If she had encountered any of these older books or plays, she would have assumed her book was mistaken since it's 'common knowledge' that Caillaux was convicted.

Intuitions may vary on this, but in Variant H, I don't think Charlotte knows that *p*. If that's right, then whether Charlotte knows that *p* is sensitive not just to the evidence she has, but to the evidence that is all around her. If she's swimming in a sea of evidence against *p*, and by the sheerest luck has not run into it, the evidence she does not have can block knowledge that *p*.

The previous example relied on the possibility of counter-evidence being everywhere. Possibly all that matters is that the counter-evidence is in just the right *somewhere*.

Variant S

In this world, an over-zealous copy-editor makes a last minute change to the very first printing of Clark's text. Not able to believe that Caillaux was acquitted—the evidence was so conclusive—they change the word 'acquittal' to 'conviction' in the sentence describing the end of the trial. Happily, this error is quickly caught, and only the first printing of the book contains the mistake. Charlotte discovered the book in a second-hand shop, which had two copies—one from the flawed first printing and one from a later printing. She bought the later one simply because it was the first one she saw. If she had entered the history section from the other direction, she would have bought the first printing and believed that *p* was false.

Plausibly, Charlotte doesn't know that *p* because it was a matter of luck that she purchased the later printing instead of the earlier one. Her method of forming beliefs, which involves buying a seemingly authoritative history book and accepting its plausible and well-supported claims, fails in this particular instance in a nearby possible world where she obtains the other copy. This type of luck is not compatible with knowledge. In contemporary terminology, a belief forming method yields knowledge only if it is *safe*. A method is safe only if it doesn't go wrong in nearby, realistic, scenarios (Williamson, 2000). So whether one knows is sensitive to not just the evidence one has, but the evidence one could easily have had.

Safety in this sense is a tricky notion. In Variant K, it seems to me that Charlotte does know that *p*.

Variant K

Charlotte detests reading books on paper, and only ever reads on her Kindle (an electronic book-reading device). Just like in Variant S, there was an error in the first printing of Clark's book. But the Kindle version never contained this error, and in any case, Kindle versions are updated frequently so even if it had, the error would have been quickly corrected. Charlotte reads the book on her Kindle, and comes to believe that *p*.

In this case, Charlotte believes *p* on good evidence from a trustworthy source, and there is no realistic possibility where she goes wrong on this question by trusting this source. That seems to me like enough for

knowledge. I'll return to the difference between Variants S and K in Section 7.4, but first I want to look at two more cases.

Variant C

Charlotte reads Clark's book and believes p . But like in Variant B, she was sure that Caillaux would be convicted. And she still thinks it is absurd that someone would be acquitted given this evidence. Rather than responding to these conflicting pressures by withholding judgment, she responds by both believing that p is true, and believing it is false. She is just inconsistent, like so many of us are in so many ways.

It seems to me that in this case, Charlotte does not know that p . The incoherence in her beliefs on this very point undermines her claim to knowledge. With one more change, we get to the case that motivates this book.

Variant I

Charlotte reads the book, and believes that p . She is then offered a bet by a curiously benevolent deity. If she takes the bet, and p is true, she wins a dinner at her favourite bistro, *Le Temps des Cerises*. If she takes the bet, and p is false, she is cast into The Bad Place for eternity. If she declines the bet, life goes on as normal. Now she's deciding what to do.

By this stage you won't be surprised to hear that I think Variant I is just like Variant C in being a case in which Charlotte lacks knowledge. What I want to defend is something even stronger than that. In Variants C and I Charlotte lacks knowledge for just the same reason; it would be incoherent to believe p . Knowledge requires coherence and rationality, and in Variant I, if Charlotte believes p , she is either irrational or incoherent. I'll come back to this point about the relationship between Variants C and I in Section 7.3. First I want to reflect a bit on what we've seen in the earlier cases.

Most of the people who think that it is implausible that interests matter to knowledge are happy acknowledging the varieties of sensitivity that are revealed by Variants J, T, B, D, H, S, K, and C. (Or at least they acknowledge most of these; maybe they have idiosyncratic objections to including one or other kind of sensitivity.) They just think this one new kind of sensitivity is a bridge too far. It is a bit of a puzzle to me why we should think sensitivity to interests is more philosophically problematic than the other kinds of sensitivity we've seen so far. It might help to get you to share my puzzlement by starting with what looks like

a simple question. What should we call the class of factors knowledge is sensitive to which is revealed by these variants, but which does not include interests?

One option is to call them the ‘traditional’ factors. Now since discussion of, say, safety only really became widespread in the 1990s, the tradition of including it in one’s theory of knowledge is quite a new one. But I don’t mind calling new things traditional. I’m Australian, and we have great traditions like the traditional Essendon-Collingwood Anzac Day match, which also dates to the 1990s. This terminology is a bit unstable though. After all, we’ve been discussing the role of interests in epistemology since at least 2002 (Fantl and McGrath, 2002), so that’s almost long enough to be traditional as well.

Another option is to say that they are the factors that are truth-connected, or truth-relevant. But there’s no way to make sense of this notion in a way that gets at what is wanted. For one thing, it’s really not obvious that coherence constraints (like we need for Variant C) are connected to truth. For another, all Variant I suggests is that we need a principle like the following in our theory of knowledge.

Someone knows something only if their evidence is strong enough for them to rationally treat the thing as a fixed starting point in their inquiries.

On the face of it, that’s at least as truth-connected as the relatively uncontroversial requirement that knowledge be based on evidence. It just says knowledge requires strong evidence. Now, of course, it also says just how strong the evidence must be depends on what their inquiries are. Is that problematic? It might be if you think that every aspect of a requirement on knowledge is truth-relevant.

That last claim really can’t be right. Or, at least, it can’t be right unless you believe the JTB theory of knowledge. If the JTB theory is false, then any premise one might use in a Wright-style argument against IRT is bound to have counterexamples. Recall the particular way Wright argued against IRT

X didn’t (have enough evidence to) know P at t but does at t* and has exactly the same body of P-relevant evidence at t* as at t. (Wright, 2018: 368)

If evidence primarily affects justification, then similarity of evidence at t and t^* should just tell us that X is rational in believing P at both times or neither. Let's say that it's both times. Then as long as one could be in a JTB-but-not-knowledge-situation at t and a knowledge-with-the-same-evidence-situation at t^* , Wright's conjunction should be possible. Here's one way that could happen.

*Variant S**

Charlotte reads the book on her Kindle, and believes that p at t_0 . The next day, at t , she can't believe she read that p and reads the book again. It still says that p , but this is bizarre because a new version of the book that says $\neg p$ was pushed out to all Kindles. Due to a network failure, Charlotte's Kindle was the only one not to get the push. She now doesn't know that p ; this case is just like the safety cases and the Harman cases. The next day at t^* a corrected version of the book that says p is pushed out to all Kindles, including Charlotte's. Again perplexed, she triple checks, and comes to believe, and know, that p .

The ugly conjunction that IRT endorses is something that theories that are sensitive to safety considerations, or evidential availability considerations, also endorse. And the true theory is sensitive to one or other kind of these considerations.

7.3 Making Up Knowledge

All that said, I've come to think there is something right about the JTB theory. Or, as I'd prefer, the RTB theory; as in Rational True Belief.² It isn't extensional adequacy; Dharmottara refuted that 1300 years ago. But it can be expressed using the modern³ notion of grounding. Or, as I'd prefer, using the notion of a *building relation* that Bennett (2017) describes.

Consider a very abstractly described case where all of 1–4 are true.

1. S knows that p .
2. p .

2 I think it's strange to apply the notion of justification to beliefs, and much more natural to talk about rational beliefs.

3 Well, modern if you think it's not the same notion as Meister Eckhart's notion of grounding. I'm a little agnostic on that.

3. S's attitude to *p* is rational.
4. S believes that *p*.

I think that when 1 is true, it is made true by 2–4. Following Bennett, we might say that the fact expressed in 1 is built from the facts expressed in 2–4. Now to make this work, we need a notion of building (or grounding) that's contingent, since 2–4 do not collectively entail 1. Defending the coherence of such a notion in detail would make for a very different book to this one. But I'll say a few words about why I think such a notion is going to be needed.

When I say that 1 is made true by 2–4, I mean that it is metaphysically explained by 2–4. They provide a complete explanation of 1's truth. Now here's the key step. A complete explanation need not be an entailing explanation. I'll give a relatively uncontroversial example of this involving causal explanation, then suggest a different philosophical example.

It is, famously, hard to explain the origins of World War One. But without settling all the causal and explanatory issues about the war's origins, we can confidently make the following two claims.

C

Had a giant asteroid struck Sarajevo on June 27, 1914, the war would not have started when it did.

NE

It is no part of the explanation of the start of the war that no such giant asteroid struck Sarajevo on June 27, 1914.

The counterfactual claim, C, can easily be verified by thinking about the consequences of giant asteroid strikes. (See, for example, the extinction of the dinosaurs.)

The claim about explanation, NE, can be verified by thinking about how absurd the task of explanation would be if it were false. For every possible event that could have changed history, but didn't, we'd have to include its non-happening in our explanation of the war. The non-occurrence of every possible alien invasion, mass pandemic, or tulip mania that could have happened, and would have made a difference, would be part of our explanation.

So the origins of the war are sensitive to whether there was a giant asteroid strike, but the lack of a giant asteroid strike is no part of the complete explanation for why the war took place. Complete causal explanations can leave out things that are counterfactually relevant to whether the event took place. That means that they aren't entailing explanations, since if everything in the complete explanation happened, but so did an asteroid strike, the war wouldn't have taken place.

We see the same thing in commonsense morality. This is one of the key points behind Bernard Williams's "One Thought Too Many" argument (Williams, 1976). If one's child is drowning in a pool, one has a reason to dive in and rescue them. Moreover, it's a complete reason. When someone asks, "Why did you do that?", you've given them a complete reason if you say, "My child was drowning". And you should accept that answer even if you think there are cases where that would be the wrong thing to do. Set up your preferred horror story moral example where diving in to rescue the child would lead to the destruction of the world. Had that horror story been actual, it perhaps would not have been morally required to dive into the pool. But in reality, a complete explanation of why it was required was that one's child was drowning.

The same thing is true about the relationship between knowledge and interests. What one knows is always (in principle) sensitive to what one's interests are. But in cases where one knows, one's knowledge is not explained by what one's interests are. Rather, it is explained just by the factors that go into RTB, and perhaps the interplay between them.

Some of the objections to IRT might rely on conflating the concepts of building and of counterfactual dependence. In their critique of IRT, Gillian Russell and John Doris (2009) repeatedly talk about how implausible it is that a change in interests can "make" one have knowledge. Strictly speaking, I don't think a change in interests does make one have knowledge. It's true that one might have knowledge, and not have had that knowledge had one's interests been different. But it doesn't follow that facts about interests stand in a making, or building, relationship to facts about knowledge. They could be, and should be, treated as things relevant to whether facts about truth, belief and rationality suffice in the circumstances for knowledge. Those factors, and only those factors, make for knowledge. That's true whether we're talking about familiar counterexamples to the JTB (or RTB) theories, or whether we're talking about interest-relativity.

The distinction between building and counterfactual sensitivity explains part of why the verdicts of IRT can sound implausible, but it doesn't explain all of it. To defend IRT from the claim that it renders implausible verdicts, we need something more. So I'll end this chapter with an argument by Nilanjan Das that responds to this kind of objection. The argument is going to be that every plausible theory of knowledge is committed to some kind of interest-relativity, and so the intuitions that my version of IRT violates are violated by every plausible theory of knowledge. Such intuitions must be wrong, and therefore can't form the basis of a good objection.

7.4 Every Theory is Interest-Relative

Think about the difference between Variant S and Variant K.⁴ Variant S was meant to be a simple case where Charlotte does not know something because of a safety violation. Knowledge is incompatible with a certain kind of luck. To know something is to do better than make a lucky guess. Charlotte isn't guessing, but she seems to be lucky in a similar kind of way to the guesser, so she doesn't know. But in Variant K, she isn't lucky. It's no coincidence that her book said the correct thing. There is no serious possibility of her being misled on this point.

Since Charlotte knows that *p* in Variant K, but not in Variant S, knowledge is sensitive to one's preferred format for reading books. This is hardly a 'truth-relevant' feature, so knowledge isn't only relevant to truth-relevant features. Knowledge generally depends on whether one was lucky, and the factors that determine whether one was lucky on an occasion need not be truth-relevant.

The same pattern recurs in other cases. In Variant H, Charlotte lacked knowledge because of evidence around her. But imagine a variant of that variant where Charlotte recently emigrated to a country where no one ever talks about Henriette Caillaux. In the variant, Charlotte knows that *p*. So her knowledge of French history is sensitive to her emigration status. And emigration status isn't truth-relevant or truth-connected.

If knowledge is sensitive to external factors, and it isn't required that knowledge be infallible, then knowledge will be sensitive to things that

4 Though they are making somewhat different points, there is a resemblance between these cases and the cases that Gendler and Hawthorne (2005) use to raise trouble for fake barn intuitions.

are not particularly truth-relevant. Any fallibilist, externalist theory of knowledge will have to face a version of a reference class problem, in order to say whether a particular true belief was a matter of luck. In general, the factors that place one in this reference class rather than another are not truth-relevant, but they are relevant to whether one can be said to know something.

That's enough to argue against sweeping generalisations about what knowledge could or could not be sensitive to. Knowledge could be sensitive to anything, because anything could matter to which reference class one is in. Das (2016: 116) shows that we can say something stronger. Cases like these can be used to directly argue for interest-relativity, even if one rejects all the other arguments in the existing literature on IRT.

Knowledge requires not getting it right just by luck. Making that intuition precise is a lot of work, but it means at least that the following is true. If the method the person used to form their belief frequently goes wrong in their actual environment, then even on occasions that the method gets the right answer, it isn't knowledge. But what's their environment? It's not just spaces within a fixed distance from them. Rather, it's spaces that in which they could easily have ended up being. It's spaces where it's a matter of luck that they are or aren't in them. So my environment, in the relevant sense, consists of a network of college towns and universities throughout the globe, and excludes any number of places a short drive away. But should I become more interested in nearby suburbs than far away colleges, my environment would change. That is to say, environment is an interest-relative notion.

If knowledge is sensitive to what one's environment is like, and one's environment in the relevant sense is interest-relative, then knowledge is going to be interest-relative. That's what is going on with Charlotte and the Kindle. Two people can be alike in what signals they receive from the world, and alike in their immediate surroundings, but be in different environments because of their different interests. If the method they use to form beliefs on the basis of that signal has differing levels of success in different environments, then whether they have knowledge will be sensitive to which environment they are in. That will depend on any number of 'non-traditional' factors, including their interests.

Now this isn't the only way, or even the main way, that interests matter to knowledge. But it is a way. And it shows that objections that

rely on the very idea of knowledge being interest-relative must overgenerate. Unless such objections are tied to a rejection of the idea that safety or reliability or any other external factor matters to knowledge, they rule out too much.

That concludes the defence of IRT over the last three chapters. The final two chapters of the book return to setting out the view, going over two important, but technical, points. First, I argue that rational belief is not sensitive to interests in quite the same way that knowledge is. And second, I argue that evidence is interest-relative, but also in not quite the same way that knowledge is.

