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All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at https://archive.org/web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

2.1 Red or Blue?

The key argument that knowledge is interest-relative starts with a puzzle about a game. Here are the rules of the game, which I'll call the Red-Blue game.

- 1. Two sentences will be written on the board, one in red, one in blue.
- 2. The player will make two choices.
- 3. First, they will pick a colour, red or blue.
- 4. Second, they say whether the sentence in that colour is true or false.
- 5. If they are right, they win. If not, they lose.
- 6. If they win, they get \$50, and if they lose, they get nothing.

Our player is Anisa. She has been reading some medieval history, and last night was reading about the Battle of Agincourt. She was amused to see that it took place on her birthday, October 25, and in 1415, precisely six hundred years before her own birthday. The book says all these things about the Battle of Agincourt because they are actually true, and when she read the book, Anisa believed them. She believed them because she had lots of independent evidence that the book was reliable (it came from a respected author and publisher, it didn't contradict her wellgrounded background beliefs), and she was sensitive to that evidence of its reliability. These beliefs were correct; the book was reliable and accurate on this point. The Battle of Agincourt was indeed on October 25, 1415, and everything else the book says about the battle without qualification is also true. Anisa comes to know that she is playing the Red-Blue game, and that these are its rules. She does not come to know any other relevant fact about the game.¹ When the game starts, the following two sentences are written on the board, the first in red, the second in blue.

- Two plus two equals four.
- The Battle of Agincourt took place in 1415.

Anisa looks at this, thinks to herself, "Oh, my book said that the Battle of Agincourt was in 1415, so (given the rules of the game) playing Blue-True will be as good as any other play, so I'm playing Blue-True. Playing Red-True would get the same amount, since obviously two plus two is four, but I'm going to play Blue-True instead". That's what she does, and she wins the \$50.

Intuitively, Anisa's move here is irrational, because it creates a needless risk. There was a simple safe option that she should have taken, and she declined it. Now it wasn't that much money; it's \$50. To be sure, she doesn't actually lose it; she gets the answer correct. The worlds where the risk is costly are somewhat distant; they are worlds where either she has misremembered something that seems vivid, or where a book that is clearly reliable has gone wrong. Still, it's sometimes true that books, even good ones, make mistakes, and memory falters. She took a risk, one that she didn't have to take, and got no compensation for taking it. That's irrational.

I'm going to argue, at some length, that the best explanation of why it is irrational for Anisa to play Blue-True is that knowledge is interestrelative. When she was at home reading the book and just thinking about

¹ When presenting this material, some people have been puzzled about how this could be possible. It's implausible that Anisa knows nothing else about the game; if she didn't know who was putting the money up she could hardly trust that she would be paid out iff she was correct. More importantly, this extra knowledge might tell her something about the sentences. I think it helps assuage these worries to imagine this as one round of a repeated game Anisa is playing. Every round two sentences from a large stock are drawn at random to be the red and blue sentences. Anisa will play 20 such rounds, and get paid something between \$0 and \$1000 at the end, depending on how many she gets right. Why is she playing this? It could be the prize round of a game show that she was the nightly winner on. With something like this background, it's plausible that what I said in the text is true; she knows 1–6, and nothing else relevant. At least, this backstory should be enough to make it plausible that the setup is indeed possible.

medieval history, Anisa knew that the Battle of Agincourt took place in 1415. When she was playing the game, and thinking about winning as much money as possible, Anisa did not know this. When she is moved into the game situation, she loses some knowledge she previously had.

In the recent literature, arguments for and against interest-relativity to date have not focussed on examples like Anisa's, but on examples involving high-stakes choices. I'll present one example, involving a character I'll call Blaise, presently. The example involving Anisa does, however, have a handful of notable predecessors. Its structure is similar to the examples of low-cost checking that Bradley Armour-Garb (2011) discusses (though he draws contextualist conclusions from these examples, not interest-relative ones). And it is similar to some of the cases of three-way choice that Charity Anderson and John Hawthorne deploy in arguing against interest-relativity (2019a, 2019b). Still, these are outlier cases. Most of the literature has focussed on high-stakes cases. Let's have one on the table.

Last night, Blaise was reading the same book that Anisa was reading. He too was struck by the fact that the Battle of Agincourt took place on October 25, 1415. Today he is visited by a representative of the supernatural world, and offered the following bet. (Blaise knows these are the terms of the bet, and doesn't know anything else relevant.) If he declines the bet, life will go on as normal. If he accepts, one of two things will happen.

- If it is true that the Battle of Agincourt took place in 1415, an infant somewhere will receive one second's worth of pure joy, of the kind infants often get playing peek-a-boo.
- If it is false that the Battle of Agincourt took place in 1415, all of humanity will be cast into The Bad Place for all of eternity.

Blaise takes the bet. The Battle of Agincourt was in 1415, and he can't bear the thought of a lovable baby missing that second of pure joy.

Again, there is an intuition that Blaise did something horribly wrong here, and one possible explanation of this wrongness is that knowledge is interest-relative. However, the argument that the interest-relativity of knowledge is the very best explanation of what's going on is somewhat weaker in Blaise's case than in Anisa's. It's not that I don't accept the interest-relative explanation of the case; I do accept it. It's rather that plausible interest-invariant explanations of the intuitions about Blaise's case exist. Because these competing explanations exist, it's hard to argue that interest-relativity is the best explanation of why Blaise's action is wrong. Without that argument, it's hard to infer from Blaise's case that knowledge is interest-relative by inference to the best explanation. So I'll focus on Anisa, not Blaise.

This choice of focus occasionally means that this book is less connected to the existing literature than I would like. I occasionally infer what a philosopher would say about cases like Anisa's from what they have said about cases like Blaise's. I'll probably get some of those inferences wrong. But I want to set out the best argument for the interest-relativity of knowledge that I know, and that means going via the example of Anisa.

Though I am starting with an example, and with an intuition about it, I am not starting with an intuition about what is known in the example. I don't have any clear intuitions about what Anisa knows or doesn't know while playing the Red-Blue game. The intuition that matters here is that her choice of Blue-True is irrational. It's going to be a matter of inference, not intuition, that Anisa lacks knowledge.

That inference will largely be by process of elimination. In Section 2.2 I will set out four possible things we can say about Anisa, and argue that one of them must be true. (The argument won't appeal to any principles more controversial than the Law of Excluded Middle.) But all four of them, including the interest-relative view I favour, have fairly counterintuitive consequences. So something counterintuitive is true around here. This puts a limit on how we can argue. At least one instance of the argument *this is counterintuitive, so it is false* must fail. That casts doubt over all such arguments. This is a point that critics of interest-relativity haven't sufficiently acknowledged, but it also puts constraints on how one can defend interest-relativity.

When Anisa starts playing the Red-Blue game, her practical situation changes. You might think I've gone wrong in stressing Anisa's interests, not her practical situation. I've put the focus on interests for two reasons. One is that if Anisa is totally indifferent to money, then there is no rational requirement to play Red-True. We need to posit something about Anisa's interests to even get the data point that the interest-relative theory explains. The second reason, which I'll talk about more in Section 2.5, is that sometimes we can lose knowledge due to a change not in our practical situation, but our theoretical interests.

In the existing literature, views like mine are sometimes called versions of **subject-sensitive invariantism**, since they make knowledge relevant to the stakes and salient alternatives available to the subject. This is a bad name; of course whether a knowledge ascription is true is sensitive to who the subject of the ascription is. I know what I had for breakfast and you (probably) don't. The distinctive feature of theories like mine is that a particular fact about the subject's situation is relevant: their interests. That should be reflected in the name. In the past, I've called this view **interest-relative invariantism**, or IRI. For reasons I'll say more about in Section 2.7, I'm not committed to *invariantism* in this book. So in this book, I will refer to it simply as the interest-relative theory of knowledge (IRT).

2.2 Four Families

A lot of philosophers have written about cases like Anisa's and Blaise's over the last couple of decades. Relatedly, there are a huge number of theories that have been defended concerning these cases. Rather than describe them all, I'm going to start with a taxonomy of them. The taxonomy has some tricky edge cases, and it isn't always trivial to classify a philosopher from their statements about the cases. It is, nevertheless, a helpful way to start thinking about the available moves.

Our first family of theories are the **sceptical** theories. They deny that Anisa ever knew that the Battle of Agincourt was in 1415. The particular kind of sceptic I have in mind says that if someone's epistemic position is, all things considered, better with respect to q than with respect to p, that person doesn't know that p. The core idea for this sceptic, which perhaps they draw from work by Peter Unger (1975), is that knowledge is a maximal epistemic state, so any non-maximal state is not knowledge. The sceptics say that for almost any belief, Anisa's belief that two plus two is four will have higher epistemic standing than that belief, so that belief doesn't amount to knowledge.

Our second family of theories are what I'll call **epistemicist** theories. The epistemicists say that Anisa's reasoning is perfectly sound, and perhaps Blaise's is too. They both know when the Battle of Agincourt took place, so they both know that the choices they take are optimal, so they are rational in taking those choices. The intuitions to the contrary are, say the epistemicist, at best confused. There is something off about Anisa and Blaise, perhaps, but it isn't that these particular decisions are irrational.

It's not essential to epistemicism, but one natural form of epistemicism takes on board Maria Lasonen-Aarnio's point that act-level and agent-level assessments might come apart.² On this version of epistemicism, taking the bet reveals something bad about Blaise's character, and arguably manifests a vice, but the act itself is rational. It's that last claim, that actions like Blaise's are rational, that is distinctive of epistemicism.

The third family is the family of **pragmatist** theories, and this family includes the interest-relative theory that I'll defend. The pragmatists say that yesterday Anisa knew when the Battle of Agincourt was, but now she doesn't. The change in her practical situation, combined with her interest in getting more money, destroys her knowledge.

And the final family are what I'll call, a little tendentiously, the **orthodox** theories. Orthodoxy says that Anisa knew when the Battle of Agincourt was last night, since her belief satisfied every plausible criterion for testimonial knowledge. Orthodoxy also says she knows it today, since changing practical scenarios or interests like this doesn't affect knowledge. On the other hand, orthodoxy says that the actions that Anisa and Blaise take are wrong; they are both irrational, and Blaise's is immoral. Moreover, it says that they are wrong because they are risky. So knowing that what one is doing is for the best is consistent with one's action being faulted on epistemic grounds.

My reading of the literature is that a considerable majority of philosophers writing on these cases are orthodox. (Hence the name!) But I can't be entirely sure, because a lot of these philosophers are more

² See Lasonen-Aarnio (2010, 2014) for more details on her view. In *Normative Externalism*, I describe the difference between act-level and agent-level assessments as the difference between asking whether what Anisa does is rational, and whether Anisa's action manifests wisdom (Weatherson, 2019: 124–125). The best form of epistemicism, I'm suggesting, says that Anisa and Blaise are rational but unwise. This isn't Lasonen-Aarnio's terminology, but otherwise I'm largely adopting her ideas.

vocal about opposing pragmatist views than they are about supporting any particular view. There are some views that are clearly orthodox in the sense I've described, and I really think most of the people who have opposed pragmatist treatments of cases like Anisa's and Blaise's are orthodox, but it's possible more of them are sceptical or epistemicist than I've appreciated.

Calling this last family orthodox lets me conveniently label the other three families as heterodox. This lets me state what I hope to argue for in this book: the interest-relative treatment of these cases is correct; and if it isn't, then at least some pragmatist treatment is correct; and if it isn't, then at least some heterodox treatment is correct.

It's worth laying out the interest-relative case in some detail, because we can only properly assess the options holistically. Every view is going to have some very counterintuitive consequences, and we can only weigh them up when we see them all laid out. For instance, here are the claims made by each of these views:

- Sceptical theories say that when Anisa is reading her book, she doesn't gain knowledge even though the book is reliable and she believes it because of a well-supported belief in its reliability.
- Epistemicist theories say that Anisa and Blaise make rational choices, even though they take what look like absurd risks.
- Pragmatist theories say that offering someone a bet can cause them to lose knowledge and, presumably, that withdrawing that offer can cause them to get the knowledge back.
- Orthodox theories say that it is irrational to do something that one knows will get the best result simply because it might get a bad result.

I'm going to mostly focus on the orthodox theories throughout the book, and in particular I'll go into much more detail on this last point in Section 2.3.

Much of the argumentation in this book, like much of what's in this literature, will fall into one of two categories. Either it will be an attempt to sharpen one of these implausible consequences, so the view with that consequence looks even worse than it does now. Or it will be an attempt to dull one of them, by coming up with a version of the view that doesn't have quite as bad a consequence. Sometimes this latter task is sophistry in the bad sense; it's an attempt to make the implausible consequence of the theory harder to say, and so less of an apparent flaw on that ground alone. Sometimes, though, it is a valuable drawing of distinctions. That is, it is scholasticism in the good sense. It turns out that the allegedly plausible claim is ambiguous. On one disambiguation we have really good reason to believe it is true, on another the theory in question violates it, but on no disambiguation do we get a violation of something really well-supported. I hope that the work I do here to defend the interest-relative theory is more scholastic than sophistic, but I'll leave that for others to decide.

Still, if all of the theories are implausible in one way or another, shouldn't we look for an alternative? Perhaps we should, but we won't find one. At least if we define the theories carefully enough, the truth is guaranteed to be among them. Let's try placing theories by asking three yes/no questions.

- 1. Does the theory say that Anisa knew last night that the Battle of Agincourt was in 1415? If no, the theory is sceptical; if yes, go to question 2.
- 2. Does the theory say that Anisa is rational to play Blue-True? If yes, the theory is epistemicist; if no, go to question 3.
- 3. Does the theory say that Anisa still knows that the Battle of Agincourt was in 1415, at the time she chooses to play Blue-True? If no, the theory is pragmatist; if yes, the theory is orthodox.

That's it—those are your options. There are two points of clarification that matter, but I don't think they make a huge difference.

The first point of clarification is really a reminder that these are families of views. It might be that one member of the family is considerably less implausible than other members. Indeed, I've changed my mind a fair bit about what is the best kind of pragmatist theory since I first started writing on this topic. There are a lot of possible orthodox theories. Finding out the best version of these kinds of theories, especially the last two kinds, is hard work, but it is worth doing. That doesn't mean that it will lessen the implausibility of endorsing a view from that family; some of the implausibility flows directly from how one answers the three questions.

The second point of clarification is that what I've really done here is classify what the different theories say about Anisa's case. They may say different things about other cases. A theory might take an epistemicist stand on Anisa's case, but an orthodox one on Blaise's case, for example. Or it might be orthodox about Anisa, but would be epistemicist if the blue sentence was something much more secure, such as that the Battle of Hastings was in 1066. If this taxonomy is going to be complete, it needs to say something about theories that treat different cases differently. So here is the more general taxonomy I will use.

The cases I'll quantify over have the following structure. Our hero, called Hero, is given strong evidence for some truth p, and they believe it on the basis of that evidence. There are no defeaters, the belief is caused by the truth of the proposition in the right way, and in general all the conditions for knowledge that people worried about in the traditional (i.e., late 20th century) epistemological literature are met. Then they are offered a choice, where one of the options will have an optimal outcome if p, but will not be the best choice according to normal theories of decision unless the probability of p is incredibly close to one. While Hero's evidence is strong, it isn't maximally strong. Despite this, Hero takes the risky option, using the fact that p as a key part of their reasoning. Now consider the following three questions.

- In cases with this form, does the theory say that when Hero first forms the belief that *p*, they know that *p*? If the answer is that this is *generally* the case, then restrict attention to those cases where they do know that *p*, and move to question 2. Otherwise, the theory is sceptical.
- 2. In the cases that remain, is Hero rational in taking the option that is optimal iff *p*? If the answer is yes in *every* case, the theory is epistemicist. Otherwise, restrict attention to cases where this choice is irrational, and move to question 3.
- 3. In *any* of the cases that remain, does the fact that Hero was offered the choice destroy their knowledge that *p*? If yes, the theory is pragmatic. If no, the theory is orthodox.

So I'm taking epistemicism to be a very strong theory—it says that knowledge always suffices for action that is optimal given what's known, and that offers of bets never constitute a loss of knowledge. The epistemicist can allow that the offer of a bet may cause a person to 'lose their nerve', and hence their belief that p, and hence their knowledge that p. Still, if they remain confident in p, they retain knowledge that p.

Pragmatism is a very weak theory—it says sometimes the offer of a bet can constitute a loss of knowledge. The justification for defending such a weak theory is that so many philosophers are aghast at the idea that practical considerations like this could ever be relevant to knowledge. So even showing that the existential claim is true—that sometimes practical issues matter—would be a big deal.

Orthodoxy is a weak claim on one point, and a strong claim on another. It says there are some cases where knowledge does not suffice for action, though it might take these cases to be very rare. It is common in defences of orthodoxy to say that the cases are quite rare and use this fact to explain away intuitions that threaten orthodoxy. The key thing is that it says that pragmatic factors never matter—so it can be threatened by a single case like Anisa.

2.3 Against Orthodoxy

The orthodox view of cases like Blaise's is that offering him the bet does not change what he knows, but still he is irrational to take the bet. In this section, I'm going to run through a series of arguments against the orthodox view. The reason I am making so many arguments is not that I lack confidence in any one of them. Rather, it is because the orthodox view is so widespread that we need to appreciate how many strange consequences it has.

2.3.1 Moore's Paradox

Start by thinking about what the orthodox view says a rational person in Blaise's situation would do. Call this rational person Chamari. According to the orthodox view, offering someone a bet does not make them lose knowledge. So Chamari still knows when the Battle of Agincourt was fought. Chamari is rational, so despite having this knowledge, Chamari will decline the bet. Think about how Chamari might respond when you ask her to justify declining the bet.

You: When was the Battle of Agincourt? Chamari: October 25, 1415. You: If that's true, what will happen if you accept the bet? Chamari: A child will get a moment of joy. You: Is that a good thing? Chamari: Yes. You: So why didn't you take the bet? Chamari: Because it's too risky. You: Why is it risky? Chamari: Because it might lose. You: You mean the Battle of Agincourt might not have been fought in 1415. Chamari: Yes. You: So the Battle of Agincourt was fought in 1415, but it might not have been fought then? Chamari: Yes, the Battle of Agincourt was fought in 1415, but it might not have been fought then, and that's why I'm not taking the bet.

Chamari has given the best possible answer at each point. Yet she has ended up assenting to a Moore-paradoxical sentence. In particular, she has assented to a sentence of the form p, but it might be that not p. It is very widely held that sentences like this cannot be rationally assented to. Since Chamari was, by stipulation, the model for what the orthodox view thinks a rational person is, this shows that the orthodox view is false.

There are three ways out of this puzzle, and none of them seems particularly attractive.

One is to deny that there's anything wrong with where Chamari ends up. Perhaps in this case the Moore-paradoxical claim is perfectly assertable. I have some sympathy for the general idea that philosophers over-state the badness of Moore-paradoxicality (Maitra and Weatherson, 2010). Still, it does seem very unattractive to end up precisely here.

Another is to deny that the fact that Chamari knows something licences her in asserting it. I've assumed in the argument that if Chamari knows that p, she can say that p. Maybe that's too strong an assumption. The conversation, says this reply, goes off the rails at the very first line. From this perspective, it is hard to determine the point of knowledge. If knowing something isn't a strong enough reason to assert it, then it is unclear what would be.

The orthodox theorist has a couple of choices here, neither of them good. One is to say that although knowledge is not interest-relative, the epistemic standards for assertion are interest-relative. Basically, Chamari meets the epistemic standard for saying that p only if Chamari knows that p according to the (false!) interest-relative theory. At this point, given how plausible it is that knowledge is closely connected with testimony, it seems we would need an excellent reason to not simply identify knowledge with this epistemic standard. The other is to say that there is some interest-invariant standard for assertion. By running through varieties of cases like Anisa's and Blaise's, we can show that such a standard would have to be something like Cartesian certainty. So most everything we say, every single day, would be norm violating. Such a norm is not plausible.

So we get to the third way out, one that is only available to a subset of orthodox theorists. We can say that 'knows' is context-sensitive, that in Chamari's context the sentence "I know when the Battle of Agincourt was fought" is actually false, and those two facts explain what goes wrong in the conversation with Chamari. Armour-Garb (2011), who points out how much trouble non-contextualist orthodox theorists get into with these Moore-paradoxical claims, suggests a contextualist resolution of the puzzles. While this is probably the least bad way to handle the case, it's worth noting just how odd it is.

It's not immediately obvious how to get from contextualism to a resolution of the puzzle. Chamari doesn't use the verb 'to know' or any of its cognates. She does use the modal 'might', and the contextualist will presumably want to say that it is context sensitive. That doesn't look like a helpful way to solve the problem though, since her assertion that the battle might have been on a different day seems like the good part of what she says. What's problematic is the unqualified assertion about when the battle was, in the context of explaining her refusal to bet. We need some way of connecting contextualism about epistemic verbs to a claim about the inappropriateness of this assertion.

The standard move by contextualists here is to simply deny that there is a tight connection between knowledge and assertion (Cohen, 2004; DeRose, 2002). (So this is really a sophisticated form of a response I just rejected.) What they say instead is that there is a kind of meta-linguistic standard for assertion. It is epistemically responsible to say that *p* iff it

would be true to say *I know that p*. Since it would not be true for Chamari to say she knows when the Battle of Agincourt was fought, she can't responsibly say when it was fought.³

The most obvious reason to reject this line of reasoning is that it is implausible that meta-linguistic norms like this exist. Imagine we were conversing with Chamari about her reasons for declining the bet in Bengali rather than English, and at every line a contribution with the same content was made. Would the reason her first answer was inappropriate be that some English sentence would be false if uttered in her context, or that some Bengali sentence would be false? If it's an English sentence, it's very weird that English would have this normative force over conversations in Bengali. If it's Bengali, then it's odd that the standard for assertion changes from language to language.

If there were a human language that didn't have a verb for knowledge, then that last point could be made with particular force. What would the contextualists say is the standard for assertion in such a language? Somewhat surprisingly, no such language exists (Nagel, 2014). It's still somewhat interesting to think about possible languages that do allow for assertions, but do not have a verb for knowledge. Just what the contextualists would say is the standard for assertion in such a language is a rather delicate matter.

Rather than thinking about these merely possible languages, let's return to English, and end with a variant of the conversation with Chamari. Imagine that she hasn't yet been offered any bet, and indeed that when the conversation starts, we're just spending a pleasant few minutes idly chatting about medieval history.

You: When was the Battle of Agincourt?

Chamari: October 25, 1415.

You: Oh that's interesting. Because you know there's this bet that someone offered my friend Blaise, and I bet I could get them to offer it to you. If you were to accept it, and the Battle of Agincourt was in 1415, then a small child would get a moment of joy. Chamari: That's great, I should take that bet.

³ The objection I'm making here is really targeted at orthodox forms of contextualism. Other forms of contextualism are not subject to it. The kind of contextualism I will describe in Section 2.7.1, for instance, can agree with IRT about what's wrong with Chamari's utterances. For more on this kind of view, see Ichikawa (2017: §1.9).

You: Well, wait a second, I should tell you what happens if the Battle turns out to have been on any other date. [You explain what happens in some detail.]

Chamari: That's awful, I shouldn't take the bet. The Battle might not have been in 1415, and it's not worth the risk.

You: So you won't take the bet because it's too risky?

Chamari: That's right, I won't take it because it's too risky.

You: Why is it risky?

Chamari: Because it might lose.

You: You mean the Battle of Agincourt might not have been fought in 1415.

Chamari: Yes.

You: Hang on, you just said it was fought in 1415, on October 25 to be precise.

Chamari: That's true, I did say that.

You: Were you wrong to have said it?

Chamari: Probably not; it was probably right that I said it.

You: You probably knew when the battle was, but you don't now know it?

Chamari: No, I definitely didn't know when the battle was, but it was probably right to have said it was in 1415.

And you can probably see all sorts of ways of making Chamari's position sound terrible. The argument I'm giving here is a version of an argument against contextualism given by John MacFarlane (2005). He notes that contextualists have a particular problem with retraction; Chamari's position sounds much worse than it should if contextualism is right. Still, I don't want to put too much weight on how she sounds. Every position in this area ends up claiming some strange things. The very idea that the epistemic standard for assertion could be meta-linguistic, either in the version which says some English word determines the appropriateness conditions for assertions in every language, or that the appropriateness conditions change from language to language, is even more implausible than the idea that we should end up where Chamari does.

2.3.2 Super-Knowledge to the Rescue?

Let's leave Blaise and Chamari for a little and return to Anisa. The orthodox view agrees that it is irrational for Anisa to play Blue-True. So it needs to explain why this is so. IRT offers a simple explanation. If she plays Red-True, she knows she will get \$50; if she plays Blue-True,

she does not know that—though she knows she will get at most \$50. So Red-True is the weakly dominant option; she knows it won't do worse than any other option, and there is no other option that she knows won't do worse than any other option.

The orthodox theorist can't offer this explanation. They think Anisa knows that Blue-True will get \$50 as well. So what can they offer instead? There are two broad kinds of explanation that they can try. First, they might offer a structurally similar explanation to the one IRT gives, but with some other epistemic notion at its centre. So while Anisa knows that Blue-True will get \$50, she doesn't *super-know* this, in some sense. Second, they can try to explain the asymmetry between Red-True and Blue-True in probabilistic, rather than epistemic, terms. I'll discuss the first option in this subsection, and the probabilistic notion in the next subsection.

What do I mean here by *super-knows*? I mean this term to be a placeholder for any kind of relation stronger than knowledge that could play the right kind of role in explaining why it is irrational for Anisa to play Blue-True. So super-knowledge might be iterated knowledge. Anisa super-knows something iff she knows that she knows that ... she knows it. She super-knows that two plus two is four, but not that the Battle of Agincourt was in 1415. Or super-knowledge might be (rational) certainty. Anisa is (rationally) certain that two plus two is four, but not that the Battle of Agincourt was in 1415. Or it might be some other similar relation. My objection to the super-knowledge response won't be sensitive to the details of how we understand super-knowledge.

If a super-knowledge solution is going to work, it had better be that Anisa does not in fact super-know that the Battle of Agincourt was in 1415. That already rules out some versions of the super-knowledge solution. In normal versions of the case, Anisa does know that she knows the Battle of Agincourt was in 1415. She knows that she read this in a book, that the book had a lot of indicators of reliability, and (at least according to the orthodox theorist), that what she read was correct. If she was asked to sort people into whether they do or don't know that the battle was in 1415, she would (in normal versions of the case) be fairly good at doing this, and would sort herself into the group that does know.⁴ So she passes all the standard tests for knowing that she knows when the battle was.

For most versions of what super-knowledge is, it looks like in ideal cases it should be closed under conjunction. That is, Anisa super-knows a conjunction (that she is considering) iff she super-knows each of the conjuncts. I'll come back to one important exception to this, that super-knowledge is credence above a threshold, in the next subsection. For now, assume that super-knowledge is closed under conjunction in this way.

Given that assumption, the fact that Anisa doesn't super-know when the Battle of Agincourt was can't explain the asymmetry between Red-True and Blue-True. In particular, it can't explain why Anisa rationally must choose Red-True. This is because she doesn't super-know that playing Red-True will win the \$50. If super-knowledge is demanding enough that she doesn't know when the battle was, it's demanding enough that she doesn't know the rules of the game. That implies that she doesn't know that playing Red will win the \$50. She has ordinary testimonial knowledge of the rules, just like she has ordinary testimonial knowledge about the Battle of Agincourt. It's just as realistic that she has misunderstood the rules of the game as that a reliable history book has gotten a key date wrong. It's not just in evil demon situations that someone misunderstands a rule. In a very ordinary sense, she can't be completely certain that she has the rules correct. If testimony from careful historians can't generate super-knowledge, neither can testimony from game-show hosts.

In fact, her knowledge of the rules of the game, in the sense that matters, is probably weaker than her knowledge of history. It is not unknown for game shows to promise prizes, then fail to deliver them, either because of malice or incompetence. Knowledge of the game rules, in particular knowledge that she will actually get \$50 if she selects a true sentence, requires some knowledge of the future. That seems harder

⁴ To be sure, she presumably doesn't know for most people what they know about medieval history. What I'm imagining is that if she was presented with a bunch of people, asked if they know when the Battle of Agincourt was, and was allowed to say "Yes", "No", or "Don't Know", then most of the "Yes" and "No" answers would be correct, and she would say "Yes" about herself.

2. Interests

to obtain than knowledge of what happened in history. After all, she has to know that there won't be an alien invasion, or a giant asteroid, or an incompetent or malicious game organiser. (The last two being considerably more important considerations in normal cases.)

So there is no way of understanding 'super-knows' such that 1 is true and 2 is false.

- 1. Anisa super-knows that if she plays Red-True, she'll win \$50.
- Anisa does not super-know that if she plays Blue-True, she'll win \$50.

If the super-knowledge based explanation of why she should play Red-True worked, there should be some sense of super-knowledge where 1 is true and 2 is false. There isn't, so the explanation doesn't work.

The point I'm making here, that in thinking about these games we need to attend to the player's epistemic attitude towards the game itself, is not original. Dorit Ganson (2019) uses this point for a very similar purpose, and in turn quotes Robert Nozick (1981) making a similar point. I've belaboured it here because it is so easily overlooked. It is easy to take things that one is told about a situation, such as the rules of a game that are being played, as somehow fixed and inviolable. They aren't the kind of thing that can be questioned. In any realistic case, the rules will not have such an exalted practical or epistemic status—at least if one assumes that only what is super-known can be taken as fixed.

This is why I rest more weight on Anisa's case than on Blaise's. I can't appeal to your judgment about what a realistic version of Blaise's case would be like, because there are no realistic versions of cases like Blaise's. Anisa's case, on the other hand, is very easy to imagine and understand. We can ask what a realistic version of it would be like. That version would be such that the player would know what the rules of the game are, but would also know that sometimes game shows don't keep their promises, sometimes they don't describe their own games accurately, sometimes players misinterpret or misunderstand instructions, and so on. This shouldn't lead us to scepticism: Anisa knows what game she's playing. But she doesn't super-know what game she's playing, which means she doesn't super-know she'll win if she plays Red-True.

2.3.3 Rational Credences to the Rescue?

So imagine the orthodox theorist drops super-knowledge, and looks somewhere else. A natural alternative is to use credences. Assume that the probability that the rules of the game are as described is independent of the probabilities of the red and blue sentence. Assume also that Anisa must, if she is to be rational, maximise expected utility. Then we get the natural result that Anisa should pick the sentence that is more probably true.⁵ And that can explain why she must choose Red-True, which is what the orthodox theorist needed to explain.

This kind of approach doesn't really have any place for knowledge in its theory of action. One should simply maximise expected utility; since doing what one knows to be best might not maximise expected utility, we shouldn't think knowledge has any particularly special role.

There are many problems with this kind of approach. Several of these problems will be discussed elsewhere in this book at more length. I will point to where those problems are discussed rather than duplicate the discussion here. Some other problems I'll address straight away.

Like the view discussed in Section 2.3.1 that separates knowledge from assertion, separating knowledge from action leads to strange consequences. As Timothy Williamson (2005) points out, once we break apart knowledge from action in this way, it becomes hard to see the point of knowledge. It's worth pausing a bit more over the bizarreness of the claim that Blaise knows that taking the bet will work out for the best, but he shouldn't take it—because of its possible consequences.

If one excludes knowledge from having an important role in one's theory of decision, one ends up having a hard time explaining how dominance reasoning works. It is, however, a compulsory task for a theory of decision to explain how dominance reasoning works. Among other things, we need a good account of how dominance reasoning works in order to handle Newcomb's problem, and we need to handle Newcomb's problem in order to motivate, or even to state, a careful version of expected utility maximisation. That little argument was very

⁵ Strictly speaking, we need one more assumption—namely that for any unexpected way for the game to be, the probability of it being that way is independent of the truth of both the red and blue sentences. This feels like a safe assumption for the orthodox theorist to make.

compressed. I'm not going to expand upon it just yet because there will be so much more discussion of dominance reasoning throughout this book; a sketch will do for now.

Probabilistic models of reasoning and decision have their limits, and what we need to explain about the Red-Blue game goes beyond those limits. So probabilistic models can't be the full story about the Red-Blue game. To see this, imagine for a second that the Blue sentence is not about the Battle of Agincourt, but is instead a slightly more complicated arithmetic truth, like *Thirteen times seventeen equals two hundred and twenty one*, or a slightly complicated logical truth, like $\neg q \rightarrow ((p \rightarrow q) \rightarrow \neg p)$. If either of those are the blue sentence, then it is still uniquely rational to play Red-True, even though the probability of each of those sentences is one. So rational choice is more demanding than expected utility maximisation. In Sections 8.2 and 8.3 I'll go over more cases of propositions whose probability is 1, but which should be treated as uncertain even it is certain that two plus two is four. The lesson is that we can't just use expected utility maximisation to explain the Red-Blue game.

Finally, we need to understand the notion of probability that's being appealed to in this explanation. It can't be some purely subjective notion, like credence, because that couldn't explain why some decisions are rational and others aren't. If Anisa was subjectively certain that the Battle of Agincourt was in 1415, she would still be irrational to play Blue-True. It can't be some purely physical notion, like chance or frequency, because that won't even get the cases right. (What is the chance, or frequency, of the Battle of Agincourt being in 1415?) It needs to be something like evidential probability. That will run into problems in versions of the Red-Blue game where the Blue sentence is arguably (but not certainly) part of the player's evidence. I'll end my discussion of orthodoxy with a discussion of cases like these.

2.3.4 Evidential Probability

No matter which of these explanations the orthodox theorist goes for, they need a notion of evidence to support them.⁶ Let's assume that we

⁶ This subsection is based on Weatherson (2018: §2).

can find some doxastic attitude D such that Anisa can't rationally stand in D to *Play Blue-True*, and that this is why she can't rationally play Blue-True. Then we need to ask the further question, why doesn't she stand in relation D to *Play Blue-True*? And presumably the answer will be that she lacks sufficient evidence. After all, if she had optimal evidence about when the Battle of Agincourt was, she could play Blue-True.

The orthodox theorist also has to have an interest-invariant account of evidence. I guess although it's logically possible to have evidence be interest-relative, but knowledge be interest-neutral, it is very hard to see how one would motivate such a position.

Now we run into a problem. Imagine a version of the Red-Blue game where the blue sentence is something that, if known, is part of the player's evidence. If it is still irrational to play Blue-True, then any orthodox explanation that relies on evidence sensitive notions (like super-knowledge or evidential probability) will be in trouble. The aim of this subsection is to spell out why this is.

Let's imagine a new player for the Red-Blue game. Call her Parveen. She is playing the game in a restaurant near her apartment in Ann Arbor, Michigan. Just before the game starts, she notices an old friend, Rahul, across the room. Rahul is someone she knows well, and can ordinarily recognise, but she had no idea he was in town. She actually thought Rahul was living in Italy. Still, we would ordinarily say that she now knows Rahul is in town; indeed, that he is in the restaurant. As evidence for this, note that it would be perfectly acceptable for her to say to someone else, "I saw Rahul here". Now the game starts.

- 3. The red sentence is *Two plus two equals four*.
- 4. The blue sentence is *Rahul is in this restaurant*.

On the one hand, there is only one rational play for Parveen: Red-True. She hasn't seen Rahul in ages, and she thought he was in Italy. A glimpse of him across a crowded restaurant isn't enough for her to think that *Rahul is in this restaurant* is as likely as *Two plus two equals four*. She might be wrong about Rahul, so she should take the sure money and play Red-True. So playing the Red-Blue game with these sentences makes it the case that Parveen doesn't know where Rahul is. This is another case where knowledge is interest-relative, and at first glance it doesn't look very different to the other cases we've seen.

But take a second look at the story for why Parveen doesn't know where Rahul is. It can't be just that her evidence makes it certain that two plus two equals four, but not certain that Rahul is in the restaurant. At least, it can't be that unless it is not part of her evidence that Rahul is in the restaurant. If evidence is not interest-relative, then it is part of Parveen's evidence that Rahul is in the restaurant. This isn't something she infers; it is a fact about the world she simply appreciates. Ordinarily, it is a starting point for her later deliberations, such as when she deliberates about whether to walk over to another part of the restaurant to greet Rahul. That is, ordinarily it is part of her evidence.

So the orthodox theorist has a challenge. If they say that it is part of Parveen's evidence that Rahul is in the restaurant, then they can't turn around and say that the evidential probability that he is in the restaurant is insufficiently high for her to play Blue-True. After all, its evidential probability is one. If they say that it is no part of Parveen's evidence that Rahul is in the restaurant because she is playing this version of the Red-Blue game, they give up orthodoxy. So they have to say that our evidence never includes things like Rahul is in the restaurant.

This can be generalised. Take any proposition such that if the red sentence was that two plus two is four and that proposition was the content of the blue sentence, then it would be irrational to play Blue-True. Any orthodox explanation of the Red-Blue game entails that this proposition is no part of your evidence—whether you are playing the game or not. Once we strip all these propositions out of your evidence, you don't have enough evidence to rationally believe, or even rationally make probable, very much at all.

Descartes, via a very different route, walked into a version of this problem. His answer was to (implicitly) take us to be infallible observers of our own minds, and (explicitly) offer a theistic explanation for how we can know about the external world given just this psychologistic evidence. Nowadays, most people think that's wrong on both counts: we can be rationally uncertain about even our own minds, and there is no good path from purely psychological evidence to knowledge of the external world. If we side with the moderns on these questions, i.e., that we do not have infallible access to our own minds, and that there is no theistic proof of the external world, Descartes's position is intolerably sceptical. The orthodox position ends up being just as badly off.

2.4 Odds and Stakes

If orthodox views are wrong, then it is important to get clear on which heterodox view is most plausible.⁷ I'm defending a version of the pragmatic view, but it's a different version to the most prominent versions defended in the literature. The difference can be most readily seen by looking at the class of cases that have motivated pragmatic views.

The cases involve a subject making a practical decision. The subject has a safe choice, which has a guaranteed return of *S*. They also have a risky choice. If things go well, the return of the risky choice is S + G, so they will gain G from taking the risk. If things go badly, the return of the risky choice is *S* - *L*, so they will lose *L* from taking the risk. What it takes for things to go well is that a particular proposition *p* is true. All of this is known by the subject facing the choice. It's also true (but not uncontroversially known by the subject) that they satisfy all the conditions for knowing p that would have been endorsed by a wellinformed epistemologist circa 1997. (That is, by a proponent of the traditional view.) So *p* is true, and things won't go badly for them if they take the risk. Still, in a lot of these cases, there is a strong intuition that they should not take the bet, and as I've just been arguing, that is hard to square with the idea that they know that *p*. So assuming the traditional view is right about the subject as they were before facing the practical choice, having this choice in front of them causes them to lose knowledge that *p*.

But what is it about these choices that triggers a loss of knowledge? There is a familiar answer to this, one explicitly endorsed by Hawthorne (2004) and Jason Stanley (2005). It is that they are facing a 'high-stakes' choice. Now what it is for a choice to be high-stakes is never made entirely clear, and Anderson and Hawthorne (2019a) show that it is hard to provide an adequate definition in full generality. In the simple cases described in the previous paragraph, however, it is easy enough to say what a high-stakes case is. It just means that *L* is large. So one gets the suggestion that practical factors kick in when faced with a case where there is a chance of a large loss.

⁷ This section is based on Weatherson (2016a: §3).

The version of IRT defended in this book does not care about whether a subject faces a high-stakes bet. Instead, it says that L matters, but only indirectly. What is (typically) true in these cases is that the subject should maximise expected utility relative to their evidence.⁸ And taking the risky choice maximises expected utility only if this equation is true.

$$\Pr(p) / (1 - \Pr(p)) > L / G$$

The left-hand side expresses the odds that p is true. The right-hand side expresses how high those odds have to be before the risk is worth taking. If the equation fails to hold, then the risk is not worth taking. If the risk is not worth taking, then the subject doesn't know that p.

Since the numerator of the right-hand side is L, then one way to destroy knowledge that p is to present the subject with a situation where L is very high. It isn't, however, the only way. Since the denominator of the right-hand side is G, another way to destroy knowledge that p is to present the subject with a situation where G is very low.

In effect, we've seen such a situation with Anisa. To make the parallel to Anisa's case even clearer, consider the following case, involving a character I'll call Darja. Darja has been reading books about World War One, and yesterday read that Franz Ferdinand was assassinated on St Vitus's Day, June 28, 1914. She is now offered a chance to play a slightly unusual quiz game. She has to answer the question *What was the date of Franz Ferdinand's assassination?* If she gets it right, she wins \$50. If she gets it wrong, she wins nothing. Here's what is strange about the game. She is allowed to Google the answer before answering. So here are the two live options for Darja. In the table, and in what follows, *p* is the proposition that Franz Ferdinand was indeed assassinated on June 28, 1914.

⁸ This simplifies the relationship between rational choice and expected utility maximisation. Later in the book I'll have to be much more careful about this relationship. See Chapter 6 for many more details.

Table 2.1 Darja's choice between answering the question, and checking Google.

	p	$\neg p$
Say "June 28, 1914"	50	0
Google the answer	\$50 - ε	\$50 - ε

If Darja has her phone near her, and has cheap easy access to Google, then ε might be really low. In that case she should take the safe option; it's the one that maximises expected utility. That means she doesn't know that p, even if she remembers reading it in a book that is actually reliable. Facing a long-odds bet can cause knowledge loss, even in low-stakes situations.

So I'm committed to the view that Darja loses knowledge in her relatively low-stakes situation, and indeed I think that's true. That's not because I have any kind of intuition that she loses knowledge. I don't have any clear intuition about her case, and I'm certainly not taking any intuition about the case as a premise. What I am taking as a premise is that Darja should Google the answer in cases like this one; doing otherwise is taking a bad risk. The best explanation of why this is a bad risk is that she doesn't know when Franz Ferdinand was assassinated. So practical interests can matter even in relatively low-stakes cases.

I'm not the first to focus on these long-odds/low-stakes cases. Jessica Brown (2008: 176) notes that these cases raise problems for the stakescentric version of IRT. Anderson and Hawthorne (2019a) argue that once we get beyond the simple two-state/two-option choices, it isn't at all easy to say what situations are and are not high-stakes choices. These cases are not problems for the version of IRT that I defend, since this version gives no role to stakes.

2.5 Theoretical Interests Matter

When explaining why I called my theory IRT, one of the reasons I gave was that I wanted theoretical, and not just practical, interests to matter to knowledge.⁹ This is also something of a break with the existing

⁹ This section is based on Weatherson (2017: §4).

2. Interests

literature. After all, Stanley's book on interest-relative epistemology is called *Knowledge and Practical Interests*. He defends a theory on which what an agent knows depends on the practical questions they face. There are strong reasons to think that theoretical reasons matter as well.

In Section 2.4, I suggested that someone knows that p only if the rational choice to make would also be rational given p. That is, someone knows that p only if the answer to the question *What should I do?* is the same unconditionally as it is conditional on p. My preferred version of IRT generalises this approach. Someone knows that p only if the rational answer to a question she is interested in is the same unconditionally as it is conditional on p. My preferred version of IRT generalises this approach. Someone knows that p only if the rational answer to a question she is interested in is the same unconditionally as it is conditional on p. Interests matter because they determine just what it is for the person to be interested in a question. Are the questions, in this sense, always practical questions, or do they also include theoretical questions? There are two primary motivations for allowing theoretical interests as well as practical interests to matter.

The first comes from what Jeremy Fantl and Matthew McGrath call the Unity Thesis (Fantl and McGrath, 2009: 73–76). They argue that whether p is a reason for someone is independent of whether they are engaged in practical or theoretical deliberation. The intuition supporting this is quite clear. Consider two people with the same background thinking about the question *What to do in situation S*. One of them is in *S*, the other is just thinking about it as an idle fantasy. Any reasoning one can properly do, the other can properly do. Since one is facing a theoretical question, and the other a practical question, the difference between theoretical and practical questions can't be relevant.

Let's make that a little less abstract. Imagine Anisa is not actually faced with the choice between Red-True, Blue-True, Red-False and Blue-False with these particular red and blue sentences. In fact, she has no practical decision to make that turns on the date of the Battle of Agincourt. Instead, she is idly musing over what she would do if she were playing that game. (Perhaps because she is reading this book.) If she knows when the battle was, then she should be indifferent between Red-True and Blue-True. After all, she knows they will both win \$50. Intuitively she should think Red-True is preferable, both in the abstract setting and when she's actually making the decision. This seems to be the totally general case.

The general lesson is that if whether one can take p for granted is relevant to the choice between A and B, it is similarly relevant to the theoretical question of whether one would choose A or B, given a choice. Since those questions should receive the same answer, if p can't be known while making the practical deliberation between A and B, it can't be known while musing on whether A or B is more choice-worthy.

There is a second reason for including theoretical interests in what's relevant to knowledge. There is something odd about reasoning from the premise that the probability of p is precisely x, to the conclusion that p, in any case where x < 1. It is a little hard to say, though, why this is problematic. We often take ourselves to know things on grounds that we would admit, if pushed, are probabilistic. The version of IRT that includes theoretical interests explains this oddity. If we are consciously thinking about whether the probability of p is x, then that's a relevant question to us. Conditional on p, the answer to that question is clearly no, since conditional on p, the probability of p is 1. So anyone who is thinking about the precise probability of p, and not thinking it is 1, is not in a position to know p. That's why it is wrong, when thinking about p's probability, to infer p from its high probability.

Putting the ideas so far together, we get the following picture of how interests matter. Someone knows that p only if the evidential probability of p is close enough to certainty for all the purposes that are relevant, given their theoretical and practical interests. Assuming the background theory of knowledge is non-sceptical, this will entail that interests matter.

2.6 Global Interest-Relativity

IRT was introduced as a thesis about knowledge. I'm going to argue in Chapter 8 that it also extends to rational belief. We need not stop there. At the extreme, we could argue that every epistemologically interesting notion is interest-relative. Doing so gives us a global version of IRT. That is what I'm going to defend here.

Stanley (2005) comes close to defending a global version. He notes that if one has both IRT, and a "knowledge first" epistemology (Williamson, 2000), then one is a long way towards global IRT. Even if

one doesn't accept the whole knowledge first package, but just accepts the thesis that evidence is all and only what one knows, then one is a long way towards globalism. After all, if evidence is interest-relative, then probability, justification, rationality, and evidential support are interest-relative too.

That's close to the path I'll take to global IRT, but not exactly it. In Chapter 9 I'm going to argue that evidence is indeed interest-relative, and so all those other notions are interest-relative too. That's not because I equate knowledge and evidence. The version of IRT I defend implies that evidence is a subset of knowledge, and which subset it is turns out to be interest-relative.

There is a deep puzzle here for IRT. On the one hand, the arguments for IRT look like they will generalise to arguments for the interest-relativity of evidence.¹⁰ On the other hand, the simplest explanation of cases like Anisa's presupposes that we can identify Anisa's evidence independent of her interests. That simple explanation says that Anisa shouldn't play Blue-True because the evidential probability of the blue sentence being true is lower than the evidential probability of the red sentence being true. Since she can't rationally play Blue-True, it follows that she mustn't know that the blue sentence is true. If evidence is identified independently, this looks like it might generalise into a nice story about when changes of interests lead to changes of knowledge. The story looks much less nice if evidence is also interest-relative, and it is.

The aim of Chapter 9 is to tell a story that avoids the worst of these problems. In the story I'll tell, evidence is indeed interest-relative, so we can't tell a simple story about precisely when changes in interests will lead to changes in knowledge. Still, it will be true that people lose knowledge when the evidential probability of a proposition is no longer high enough for them to take it for granted with respect to every question they are interested in.

¹⁰ I was first convinced of this by conversations with Tom Donaldson some years back. The earlier example of Parveen in the restaurant grew out of these conversations.

2.7 Neutrality

This book defends, at some length, the idea that knowledge is interestrelative. I am, however, staying neutral on a number of other topics in the vicinity.

2.7.1 Neutrality about Contextualism

Most notably, I'm not taking any stand on whether contextualist theories of knowledge are true or false. If you think that contextualism is true, then what I'm defending is that the view that 'knowledge' picks out in this context, and in most other contexts, is interest-relative.

Contextualist theories of knowledge have a lot in common with interest-relative theories. The kind of cases that motivate the interestrelative theories, cases like Anisa's and Blaise's, also motivate contextualism. They might even be seen as competitors, since they are offering rival explanations of similar phenomena. They are not, however, strictly inconsistent. Consider principles A and B below.

- A. A's utterance that *B* knows that *p* is true only if for any question *Q*? in which A is interested, the rational answer for B to give is the same unconditionally as it is conditional on *p*.
- B. A's utterance that *B* knows that *p* is true only if for any question *Q*? in which B is interested, the rational answer for B to give is the same unconditionally as it is conditional on *p*.

I endorse principle B, and that's why I endorse an interest-relative theory of knowledge. If I endorsed principle A, then I would be (more or less) committed to a contextualist theory of knowledge. And principle A is not inconsistent with principle B.¹¹

It isn't hard to see why cases like Anisa and Blaise can move one to endorse principle A, and hence contextualism. It would be very odd for Anisa to say "This morning, I knew the Battle of Agincourt was in 1415." That's odd because she can't now take it as given that the Battle of Agincourt was in 1415, and in some sense she wasn't in any better

¹¹ There is a technical difficulty in how to understand one person answering an infinitival question that another person is asking themselves. The points I'm making in this section aren't sensitive to this level of technical detail.

or worse evidential position this morning with respect to the date of the battle. Perhaps, and this is the key point, it would even be false for Anisa to say this now. The contextualist, especially the contextualist who endorses principle A, has a good explanation for why that's false. The interest-relative theorist doesn't have anything to say about that. Personally I think it's not obvious whether this would be false for Anisa to say, or merely inappropriate, and even if it is false, there may be decent explanations of this that are not contextualist. (For instance, maybe knowledge is sensitive to what interests one will have. Or maybe some kind of relativist theory is true.) But there is clearly an argument for contextualism here, and it isn't one that I'm going to endorse or reject.

One reason I'm not rejecting contextualism is that I'm not sure really what it is. Here's a theory about 'knows' that I think is interesting, and I don't know whether it is contextualist. The word 'knows' is polysemous. It has three possible meanings. One of them is something like Cartesian certainty. In this sense, most knowledge claims are false. Another is something like information possession. In this sense, my car might know lots of things, since its systems do quite reliably store a lot of information. Finally, there is a moderate sense, which is what we most commonly use. The difference between the three might even be marked phonologically; the Cartesian sense is often somewhat drawn out or otherwise emphasised. Is this contextualist? I don't know. Sort of, I guess. It agrees with the standard contextualist account of the appeal of scepticism. On the other hand, it denies that 'knows' has the kind of continuous variation that is typical in comparative adjectives like 'rich'. Since I think this kind of polysemy theory might be true, and (independently) that it might be contextualist, I'm not in a position to deny contextualism.

2.7.2 Other Aspects of Neutrality

As I've already noted, I'm making heavy use of the principle that Brown calls K-Suff. I'm going to defend that at much greater length in what follows. What I'm not defending is the converse of that principle, what she calls K-Nec.

K-Nec

An agent can properly use *p* as a reason for action only if she knows that *p*.

The existing arguments for and against K-Nec are intricate and interesting, and I don't have anything useful to add to them. All I will note is that the argument of this chapter doesn't rely on K-Nec, and I'm mostly going to set it aside.

I'm obviously not going to offer anything like a full theory of knowledge. I am just defending a particular necessary condition on knowledge. That condition entails that knowledge is interest-relative given some common-sense assumptions about how widespread knowledge is.

I will be making one claim about how interests typically enter into the theory of knowledge. I'll argue that there is a certain kind of defeater. A person only knows that p if the belief that p coheres in the right way with the rest of their attitudes. What's 'the right way'? That, I argue, is interest-relative. In particular, some kinds of incoherence are compatible with knowledge if the incoherence concerns questions that are not interesting.

So the impact of interests is (typically) very indirect. Even if the other conditions for knowledge are satisfied, someone might fail to know something because it doesn't cohere well with the rest of their beliefs. What turns out to be most important here is an exception to this exception clause. Incoherence with respect to uninteresting questions is compatible with knowledge.

This is going to matter because it affects how we think about what happens when interests change. It is odd to think that a change in interests could make one know something. It isn't as odd to think that a change in interests could block or defeat something that was potentially going to block or defeat an otherwise well-supported belief from being knowledge. This is something I will return to repeatedly in Chapter 7.