

A full-page background image showing a sunset over the Golden Gate Bridge. The sky is a mix of deep purple, pink, and orange, with wispy clouds. The bridge is silhouetted against the bright orange horizon. The water in the foreground reflects the colors of the sky.

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All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at 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

1. Overture

The core thesis of this book is that what a person knows is sensitive to what their interests are, and in particular to what inquiries they are engaged in. The thesis is designed to resolve a puzzle about the nature of inquiry. Inquiry has to start somewhere, and a natural place to start is with what one knows. If one is planning a meal for friends, and choosing what to make, it's natural to start with what one knows about what ingredients are on hand or easily available, what the friends like, what dietary preferences and restrictions they have, and so on.

Now we face a puzzle. Either we identify knowledge with absolute certainty or we do not. If we do, then inquiry can barely get started. If one knows anything with absolute certainty, then it is at most trivialities like instances of the law of identity. That won't be enough to get going on planning dinner. So let's say we do not identify knowledge with absolute certainty, and instead pick some particular level of certainty below that. Then there will be propositions that are more certain than that threshold, but which one should not use in this particular inquiry. For instance, there will be cases where one's evidence that a particular friend is not allergic to peanuts is just above that threshold, but given the potentially lethal consequences of getting it wrong, this isn't something that should be taken as a starting point in inquiry.

The solution is to identify knowledge with a level of certainty which varies with the nature of the inquiry. In particular, it varies both with how important it is to get the inquiry right (very important in the case of the allergy), and with how hard it would be to get further information relevant to the inquiry.

I'm not the first to defend such a view; there is a thriving literature on interest-relative theories of knowledge like the one I'm defending here. But for a long time it was a remarkably curious literature. Interest-relative theories were discussed everywhere and endorsed virtually

nowhere. It's possible things changed around the time of the COVID-19 pandemic; after 2020 there were more positive discussions of interest-relativity.¹ Before then, interest-relative theories had a ratio of discussion to endorsement that philosophy hadn't seen since David Lewis put concrete modal realism on the agenda.

The terminology that is used to describe the debate about interest-relativity is striking. The interest-relative view is usually opposed to the 'purist' or 'traditionalist' view. I'm not going to dive into the literature on which views get described as 'pure' or 'impure', but I wanted to pause a bit over 'tradition'. This is a particularly curious choice of word, and I think its curiosity is related to the strange shape of the literature around interest-relativity.

The recent literature on interest-relativity was kick-started by three works in the early 2000s. First was Jeremy Fantl and Matthew McGrath's paper "Evidence, Pragmatics and Justification", published in *The Philosophical Review* in 2002. Then came two books from Oxford University Press: *Knowledge and Lotteries* by John Hawthorne in 2003, and *Knowledge and Practical Interests* by Jason Stanley in 2004. Now these works are, by standards of recent epistemology, from quite a long time ago. That is to say, two decades is a long time in epistemology. Compare, for instance, the literature on the idea that safety is central to the theory of knowledge. The idea that safety is important plays a crucial role in a series of works from the late 1990s and early 2000s by Lewis, Timothy Williamson, Ernest Sosa, and Duncan Pritchard.² And it became a central feature of a lot of epistemological theorising very quickly. But safety-relative epistemology is really only a few years older than interest-relative epistemology. So why is one of these traditional and the other not?

One possible answer is that while safety was a new idea, it struck epistemologists as similar to older ideas. Safety looks a lot like the sensitivity condition that Robert Nozick (1981) had argued plays a

1 See, for instance: Kim (2023), Gao (2023), Schmidt (forthcoming), Steglich-Petersen (2024), Wu (forthcoming), and Ye (2024). That's about as many people defending interest-relative theories in one year as defended them for the first 15 years since they were introduced in Fantl and McGrath (2002).

2 I'll discuss safety in more detail later. For now, a rough definition will suffice. A person's belief is considered safe if they couldn't easily have been wrong in forming that belief. According to safety-relative epistemology, only safe beliefs amount to knowledge, and this plays an important role in explaining knowledge.

central role in the theory of knowledge. Sosa (1999) plays up this similarity, framing safety as a kind of converse of sensitivity. And safety looks like a kind of reliability condition, so it is continuous with 20th-century work on reliabilism. So while safety theories are new, they have things that look like precursors. But to a lot of epistemologists, interest-relative theories seemed novel. It wasn't just that they offered a new account of what affects knowledge; it was that they offered a view that came out of nowhere.

If that was the impression that epistemologists had, it was mistaken. There are precursors to contemporary interest-relative views, and looking at them is helpful for thinking about why one might want to endorse an interest-relative view. I'm going to focus on two of these precursors, one from Hellenistic philosophy and one from medieval philosophy.³

Philo of Larissa lived from around 159 BCE to around 83 BCE, and was the last sceptical head of Plato's Academy.⁴ He held a number of views over his life, but the one that's important here is his "mitigated scepticism" (Brittain and Osorio: §3.3). The sceptics faced a challenge: if no one knows anything, and indeed no one should believe anything, then it seems rational action is impossible. But surely some acts are rational, or at least more rational than other acts. What can be done?

Philo's response is to say that while it is true that nothing can be known, it can be rational to assent to certain "persuasive impressions" (Brittain and Osorio: §3.3). Action that is based in the right way on an impression that is really persuasive (and not just one that actually persuades) can be rational. Moreover, says Philo, how much evidence one needs to be properly persuaded can vary with differences in what's at stake with the action. As Peter Adamson puts it,

-
- 3 A quick note on sources. This is not at all a work of historical scholarship, and I'm not in a position to write such a work. So everything I cite here is going to be a contemporary secondary (or tertiary) source. I do hope in the future there will be more work which looks at the relationship between these historically important figures and contemporary views, but that work will have to be done by someone with a different skill set to mine.
 - 4 My main sources here are the Stanford Encyclopedia of Philosophy entry on Philo (Brittain and Osorio, 2021) and the chapters on scepticism in Peter Adamson's book on Hellenistic philosophy (Adamson, 2015). I'm particularly drawing on section 3.3 of the SEP entry, and chapters 16 and 17 of Adamson's book.

Like Arcesilaus, Philo suggests that these impressions will be used as a practical guide by the Skeptic. But he went further, observing that the standards we use will differ depending on how high the stakes are. In the normal course of affairs one bit of evidence will suffice. For instance, if I'm looking for the giraffes, I'll just ask another zoo-visitor and follow their directions. But what if it is really important—if, say, I need to be at the giraffe enclosure in five minutes to pay a ransom to the giraffe-nappers who are demanding £1 million for the safe return of Hiawatha, who just happens to be my favorite giraffe? Then I will want to make extra sure. (Adamson, 2015: 112)

Now Philo (probably) doesn't move from this to an interest-relative theory of knowledge. But look how close he gets. He thinks that the norm of belief, or at least the norm of the thing that plays the same role in his philosophical system as belief plays in ours, is interest-relative. All you have to add to get an interest-relative theory of knowledge is that knowledge is the norm of (the thing that plays the functional role of) belief.

Jumping ahead a millennium and a half, our next stop is with the epistemology of medieval philosopher Jean Buridan. I'm going to draw extensively here on Robert Pasnau's discussion of medieval epistemology in his *After Certainty*. Pasnau credits Buridan with introducing "what would become the canonical three-level distinction between absolute, natural, and moral certainty" (Pasnau, 2017: 32). The last of these, "moral certainty", is the most important one here. This isn't quite Buridan's phrase—he talks about moral evidentness—but he seems to be the causal origin of the introduction of the phrase "moral certainty" (or its equivalent in other languages) into western European discourse. And it's particularly interesting to the story here to see what kind of problem this notion is meant to solve.

There is still another, weaker evidentness, which suffices for acting well morally. This goes as follows: if someone, having seen and investigated all the attendant circumstances that one can investigate with diligence, judges in accord with the demands of such circumstances, then that judgment will be evident with an evidentness sufficient for acting well morally—even if that judgment were false on account of invincible ignorance concerning some circumstance. For instance, it would be possible for a judge to act well and meritoriously by hanging an innocent man because through testimony and other documents it sufficiently appeared to him in accord with his duty that that good man was a bad murderer. (Buridan, as quoted in Pasnau, 2017: 34)

Note in particular the phrase “the demands of such circumstances”. Buridan’s notion here is clearly interest-relative. What it takes to properly judge a defendant guilty of murder is considerably more than what it takes to judge that someone broke a promise. The difference between the misdeeds, while in the first instance a moral difference, matters to the applicability of this epistemic concept.

Now Buridan does not have an interest-relative account of knowledge. After all, the very example he uses to introduce this interest-relative concept is one where the belief is false, and hence not knowledge. This would change over time. Eventually John Wilkins, writing in the 17th century, would take moral certainty to be the standard for knowledge (Pasnau, 2017: 218). Wilkins is important to the history of science as one of the founders of the Royal Society. And he is important to the history of epistemology because he starts the tradition of centring epistemology around attainable norms. Here is how Pasnau puts the point.

Wilkins in particular, in his small way, takes what can retrospectively be seen as a decisive step, because he both rejects the principle of proportionality in favor of a broad scope for absolute belief and identifies the whole range of such belief with knowledge. For, even as he continues to associate knowledge with certainty, he allows that mere moral certainty is good enough, treating mathematical, physical, and moral as three different kinds of knowledge and thus locating the threshold for knowledge not at intellectual compulsion but at the absence of reasonable doubt: “that kind of assent which does arise from such plain and clear evidence as does not admit of any reasonable cause of doubting is called knowledge or certainty.” (Pasnau, 2017: 43)

The “principle of proportionality” here is the idea that the better one’s evidence for a proposition is, the stronger one’s belief in that proposition should be. What’s distinctive in Wilkins is that he thinks one can have absolute belief in a mere moral certainty. This violates proportionality because if one’s belief is a mere moral certainty, then the evidence for it could be improved. But since it is an absolute belief, the belief couldn’t get stronger.

What’s distinctive about Wilkins is not the use of moral certainty in epistemology. That’s there in Buridan three hundred years earlier. What’s distinctive is the central role he gives it. And as Pasnau reads the situation, the approach taken by Wilkins becomes orthodox for the next three hundred years. The alternative option, one that Pasnau prefers,

is to focus on what the epistemological ideal is, and on how close we can get to attaining that ideal. You can read at least some contemporary probabilists as working in the tradition—one that was common before Wilkins—of thinking that only maximally supported beliefs get the maximal level of belief. But this is definitely not the mainstream view for the last few centuries. The mainstream view is that there are these important, absolute, concepts that can be attained even though one's evidential situation could in principle improve further. And these concepts are closely tied to knowledge. And, most strikingly, the one that is mostly tied to knowledge is originally introduced as an interest-relative concept.

In both Philo of Larissa, and in the tradition that runs from Buridan to Wilkins and beyond, interest-relative epistemic concepts play central roles. There is no figure here who literally endorses every aspect of the contemporary interest-relative view. But the precursors are there. Indeed, they are there at some of the earliest sightings of what we might, in current terminology, call fallible epistemologies. If anything, I suspect the idea that an epistemology can be fallibilist and interest-invariant is the more recent innovation. Rather than dive too deeply into those historical waters, let's turn to a connection between Buridan's epistemology and (a particular strand in) Indian epistemology: the place of action theory in epistemology. What worries Buridan is whether a certain action, hanging an innocent man, can be given an epistemological defence. Buridan isn't the first philosopher to see a tight connection between epistemology and action theory.

The 5th century philosopher Vātsyāyana is known for his commentary on the 1st or 2nd century Nyāya-sūtra.⁵ In this commentary he offers a number of anti-sceptical arguments. This one is most interesting to the story here.

For Vātsyāyana, the purpose of knowledge is indeed crucially important. He begins his commentary by saying that knowledge is needed in order to secure any desired objective (artha). Each of us exerts effort only for the sake of achieving such an objective. Here one might think of an idea we encountered in Mīmāṃsā, that it is a sacrificer's desire that makes a ritual incumbent upon the sacrificer.

5 My source for everything here is Peter Adamson's and Jonardon Ganeri's *Classical Indian Philosophy* (Adamson and Ganeri, 2020).

No desire, no action. Now Vātsyāyana adds: no knowledge, no result! After all, how can you get what you want when you literally don't know what you're doing? Vātsyāyana invokes the point again later on, when he responds to the standard skeptical argument that any means of knowledge must be ratified by some further means of knowledge, leading to a regress. Thus, the skeptic is suggesting, we cannot trust a *pramāṇa* like perception unless some further perception tells us that it is trustworthy. No, replies Vātsyāyana. If this were true then "the activities of practical life" would be impossible, since the only way we ever achieve anything that we want is by knowing how to get it. This applies to mundane goals like wealth and pleasure, and to more exalted goals too. Nyāya competes with the Buddhists not only on the epistemological front, by refuting skeptical arguments like the one just mentioned, but also on what we may, with apologies to Monty Python, call the liberation front. The elimination of suffering, promised by Buddhists and Naiyāyikas alike, is one more objective that can be achieved through knowledge and through knowledge alone. (Adamson and Ganeri, 2020: 170)

More bluntly, the argument is that some actions are rational, only actions based on knowledge are rational, and so we have some knowledge, contra scepticism. Unlike Vātsyāyana I'm not in the business of arguing against scepticism. But this is an excellent anti-sceptical argument. That's not just because it's sound, and persuasive, though it's both. It's because it derives anti-sceptical conclusions from the practical nature of knowledge. It grounds the anti-scepticism where it should be grounded, in the practical nature of knowledge.

The Nyāya philosophers, like Vātsyāyana, are relevant to this story for another reason. As well as closely connecting knowledge with action, they connect it closely with inquiry. And this book, like many contemporary philosophers, takes the same approach. Jane Friedman (2019a, 2019b, 2020, 2024b) has developed a detailed account of what inquiry is and how it relates to epistemology. Elise Woodard (2020) and Arienne Falbo (2021) have some persuasive criticisms of particular details of Friedman's views, but enough of the picture survives, and indeed is developed by both Woodard and Falbo, to be useful in theorising about knowledge. Guido Melchior (2019) has developed a detailed account of a special kind of inquiry, namely checking, and some of what he says about checking is very useful in resolving some tensions in the interest-relative picture.

Knowledge seems like it should be related to inquiry. But just what is the relationship? An inquiry, like any action, has a beginning, a middle, and an end. And it helps to think about ways in which knowledge can play a role at each of these three stages. In particular, the following three theses suggest ways in which knowledge plays a role at each stage in turn.

1. Inquiry should start with knowledge.
2. Inquiry should only be into things one does not know.
3. Inquiry should aim at knowledge.

All three of these are plausible, but I'm ultimately only going to accept 1. I'm going to accept a fairly strong form of it. On the version I accept, only knowledge is appropriate as a starting point for inquiry, and any knowledge could (in principle) be appropriate as a starting point. The latter claim has to be qualified in some ways—it isn't appropriate to start an inquiry into where the cat is with one's knowledge about early Roman history. But if, while inquiring into where the cat is, one knows which year Hannibal crossed the Alps, then that knowledge is certain enough for use in the inquiry. If it shouldn't be used, and it probably shouldn't be, that's on grounds of irrelevance, not on grounds of uncertainty.

But I'm going to reject 2 and 3. I'm disagreeing here with Friedman, whose theory of inquiry gives an important role to 2. And Woodard argues convincingly that the failure of 3 implies that 2 has to fail as well.⁶ Inquiry might aim at knowledge, but it might aim at any number of other things: at understanding, or at sensitivity, or at developing reasons that convince others. (The latter aim plays an important role in the explanation Michael Strevens (2020) offers for some striking features of contemporary science.) Since one might want to understand something one knows, or have a more sensitive belief in what one knows, or convince others of what one knows, it can make sense to inquire into what one knows.

The next chapter presents a straightforward argument for interest-relative epistemology. Before I get to that, I want to offer two motivations

6 Note that Friedman (2024b) also rejects 3, but not because she thinks inquiry aims at something else; she is sceptical of the metaphor of aiming in this context. Note also that Falbo and Melchior developed similar arguments to Woodard's.

for the view. You could try to turn either of these motivations into a nice, clean, premise-conclusion argument for interest-relativity. I haven't done that because in both of these cases, the premise-conclusion format obscures more than it enlightens. The first motivation comes from the practical nature of belief, and the second from the thought that knowledge is a natural kind.

Think back to the problem facing Philo of Larissa. He wants to be a sceptic, so nothing is known. He also wants to be able to act in the world. Action requires a picture of what reality is like. So we need some mental state that aims to fit the world, and which can guide action. Once we have that state, you might well think that it's just belief. Hugo Mercier (2020) argues that people do not believe as many conspiracy theories as they say they do; these apparent endorsements he argues are moves in a complicated signaling game. His evidence that they don't actually believe the conspiracy theories is that they act nothing like how they would act were the theories true. Whether or not the details of Mercier's argument are right, the form of it seems right. Apparent belief that is out of sync with action is not really belief at all.

If belief is practical, the norms for belief should be practical too. This isn't a logically necessary conditional; it is easy to describe cases where we have non-practical norms for an essentially practical state. Still, you should expect that the norms of a practical state are typically practical. So you should expect that epistemology, the study of norms for belief, will be shot through with practical considerations. That's what the interest-relative theorist says is in fact the case.

The second motivation comes from reflection on what we're trying to do in epistemology, and how it relates to the importance of knowledge. I mentioned earlier that Pasnau regards the turn epistemology took after Wilkins, where a central focus is on clarifying sub-optimal notions like knowledge, to be a mistake. (By 'sub-optimal' here, I mean merely that they are standards one can meet while also being in a position to improve one's doxastic position.) He thinks this is a retreat into mere lexicography, and away from what was traditionally, and correctly, viewed as the primary task of epistemology, namely clarifying the nature of the epistemic ideal. I'm working in this post-Wilkins tradition, so I probably should say some words in defence of it. This defence ends up motivating an interest-relative approach.

Firstly, even if one didn't care about threshold standards like knowledge, the right thing to do isn't to focus on the ideal. We aren't going to attain the ideal. What we can do is get better and better. But knowing what the ideal is like is often very little help in figuring out how to do better. This is a general consequence of the Theory of the Second Best (Lipsey and Lancaster, 1956–1957). Very often, being like the ideal is a way of being worse rather than better. For example, the ideal inquirer doesn't forget anything, so they don't need to take notes while reading. Nevertheless, it's a good epistemic practice to take notes while reading. So even if what you ultimately care about is doing better, not meeting thresholds, it isn't obvious that exploring the ideal is the way to get there. If our aim is epistemic improvement, we're probably better off exploring tools that fallible humans have developed for helping other fallible humans, rather than striving for an unattainable ideal.

Secondly, and more importantly, the project here is not one of lexicography. I don't particularly care how the English word 'knows' is used. The fact that a phonologically indistinguishable word is used to talk both about knowing who won last night and knowing the players on the winning team is of no relevance to the project we're engaged in here. The fact that most languages have a word that is very close to synonymous to the English word 'knows' is more relevant. That's not because it makes the lexicography important. Rather, it's because it suggests that there is an important concept that English speakers are picking out with 'knows', that French speakers are picking out with 'savoir', and so on for all the other languages in the world. It could be that all these different language groups agreed to use one of their limited stock of words for this concept, and it was a mistake in every case. But as J. L. Austin frequently reminded us, that's not the way to bet.

The concept of knowledge is, among other things, scientifically important. Throughout the social sciences, there are theories that are grounded in patterns of human behaviour. Those patterns are, usually, best explained in terms of what those humans know. Consider the (stylised) fact that in a small, open, free market, competing suppliers of a common good will usually sell goods for the same price. We could offer an explanation of this in terms of the effective demand for a supplier's goods given their price and the price of competing suppliers. The demand curve facing this individual supplier will have a striking

discontinuity; once the price goes above the price others are offering the good at, demand falls to 0.

Such an explanation will be good as far as it goes, but we can do better. We can note that there are mechanisms—in the sense of mechanism developed by Peter Machamer, Lindley Darden, and Carl Craver (2000)—that underlie this pattern of effective demand. The mechanisms are individual consumers who will change their purchasing patterns if they know that someone else is selling the same good more cheaply. Mechanisms, in this sense, are things that display a consistent pattern of activity. The activities have external triggers and reliable outputs given that trigger. Here the trigger is knowledge that someone else is offering the good more cheaply, and the output is buying the good elsewhere. The crucial thing for us is that here, like in many other social science applications, the trigger needs to be stated in terms of knowledge. It can't just be that the change in prices leads to a change in behaviour; a change in price that no one knows about won't plausibly bring about any behavioural change. It can't be that the trigger is stated in terms of what is absolutely certain because no one can be absolutely certain of contingent things like the price that a supplier is charging for a good. Nor can the trigger be stated in terms of high probability. No matter how probable I think it is that supplier B is cheaper than supplier A, it might still be rational to buy from supplier A if the rest of the probability goes to possibilities where B is much more expensive.

Knowledge alone seems to do the trick. The generalisation that people buy from suppliers they know to be cheaper seems both to be true and to rationalise their purchasing behaviour. What's important for us is that this places knowledge at the centre of our understanding of how this social arrangement works. That is going to be the general case; you just can't do social science without talking about how people behave when they come to know things.

So we have two reasons for thinking knowledge is a reasonably natural kind: there are more or less synonymous terms for it across languages, and it plays a key role in scientific explanations. Most fallibilist theories of knowledge won't make it be particularly natural. (I'll expand on this point in Section 8.4.1.) Most such theories say that to know something is to have a belief that's good enough along some dimension. So the belief must be justified enough, or safe enough, or

produced by a reliable enough mechanism. Concepts that just pick out points high enough up some or other scale are not particularly natural. We should expect that we could do better.

Some fallibilist theories, or at least theories that make knowledge 'sub-optimal' in the sense I used above, do seem to be reasonably natural. The sensitivity theory that Nozick (1981) develops, for instance, plausibly makes knowledge into a natural kind. Whether a belief would be retained were its content false is not a matter of how well the belief performs on some scale. Alternatively, one could hold that knowledge is primitively natural, not natural in virtue of its analysis or parts.⁷ That's not completely implausible; it isn't obvious that the naturalness of social kinds has to be explained by the same things that metaphysically explain why individuals fall into those kinds. Still, it would be nice to have a better explanation of why knowledge is natural.

On the view defended here, a person knows a proposition if and only if they properly take it to be settled. What one properly takes to be settled is interest-relative, hence knowledge is interest-relative. I'm not putting forward this biconditional as an analysis of knowledge, or an explanation of knowledge. It could be that the direction of explanation here runs from knowledge to proper settling. What I am claiming is that this biconditional is true, and is part of the explanation of why knowledge is a natural kind. The way to finish that explanation is to develop a theory where knowledge is interest-relative.

So those are the two big motivations for the interest-relative view: the practicality of belief and the naturalness of knowledge. Belief is a practical notion, so the norms of it should be practical. Knowledge is, at its most essential, a norm of belief. Knowledge is a natural kind, as evidenced by its cross-linguistic prevalence and its role in science. This raises a challenge, since knowledge often feels like it requires the knower do 'well enough' along one or other scale, and there is nothing particularly natural about choosing this point on the scale rather than that point. The interest-relative theory has an answer to this problem: a believer has knowledge when their evidence is good enough to properly settle the inquiries the believer is engaged in, and that's more than an arbitrary point on a scale.

7 Such a view might be inspired by the 'knowledge first' program of Williamson (2000).

While these are motivations, neither of them is strictly speaking an argument. The main argument in this book for the interest-relative theory is developed in Chapter 2. It is that in some fairly simple situations, there is a choice between four options.

1. Accepting scepticism about all contingent knowledge.
2. Denying some very simple principles connecting knowledge and action—and in particular denying that it is rational to take the action one knows to be best.
3. Denying some very strong intuitions about which actions are rational in these simple situations.
4. Saying that knowledge is interest-relative.

Since the first three options are implausible, the fourth is correct. That is, knowledge is interest-relative.

The argument does not turn on intuitions about who knows what in what situations. The only cases where the interest-relative theory disagrees with its rivals are ones where intuitions about knowledge seem to me to be very weak. For what it's worth, and it isn't worth much, I think the interest-relative theory says the more intuitive thing about most cases. Ultimately though, I don't particularly care about intuitions about knowledge, at least in these relatively borderline cases. I will spend some time defending my version of the interest-relative theory against the frequently voiced complaint that interest-relative theories get some clear cases incorrect. When the cases are indeed clear, I'll show that my version of the theory matches the intuitions, but curve-fitting around case intuitions will not be my priority.

To know something is to properly take it to be settled. There are two kinds of practical considerations that might make it improper to take something to be settled even if the evidence in favour of settling is strong enough for everyday purposes. The first is that the cost of being wrong is very high. The second is that the cost of checking whether one is wrong is very low. The previous literature on interest-relativity has primarily focussed on the first kind of reason. So the literature is replete with discussion of 'high-stakes' cases, where someone stands to lose a lot if something they have excellent evidence for turns out to be false. Knowledge is often lost in these cases. There is somewhat less discussion,

however, about how knowledge is also lost in ‘easy checking’ cases, or about how (as Mark Schroeder (2012) notes), knowledge might not be lost in cases where the stakes are high but checking is impossible. As I’ll put it in Section 2.4, what matters is not the stakes, but the odds one faces in a particular situation.

Humans engage in both practical and theoretical inquiries. For that matter, they often engage in inquiries which mix the practical and the theoretical. A lot of the focus in the literature on interest-relativity has been on how knowledge interacts with practical inquiry. Indeed the title of Stanley’s defence of an interest-relative account is *Knowledge and Practical Interests*. I don’t impose any such restriction here. If p can’t be properly taken to be settled in a purely theoretical inquiry that someone is engaged in, they don’t know that p . This has one striking implication. Let’s say the person is trying to figure out as precisely as possible the probability of p . If they can take p as settled in that inquiry, then the answer to the inquiry will be 1. Unless the correct answer to this inquiry is actually 1, it won’t be proper to take p as settled. So in general one easy way to lose knowledge that p is true is to launch an inquiry into precisely how probable p is. I’ve set this out using the ideology of probability, but this is unnecessary. Any inquiry into how well supported p is by one’s overall evidence will usually not be allowed to take p as a starting point. So engaging in that inquiry will lead to loss of knowledge. This is what is right in scepticism, and infallibilism. The Cartesian meditator does, on this view, lose knowledge in anything when they seriously reflect on how good their evidence is for it. Happily, this knowledge comes back when they return to their normal life.

In the middle of the discussion about knowledge and probability there is a little inference: from the premise that taking p as settled would lead to an incorrect answer, it follows that it is improper to take p as settled. That’s a good inference; that taking p as settled leads to a mistaken conclusion is indeed compelling evidence that it is improper to take p as settled. But it’s not the only reason that it could be improper to take p as settled. Among other things, taking p to be settled might get to the right answer for the wrong reasons. So this principle of ‘don’t take something to be settled if it will lead to the wrong answer’ might be good advice, but it isn’t a full account of when not to take something as settled. I will go over this point in much more detail in Sections 3.4 to 3.6.

What one can properly take as given is a function of one's evidence. This should be common ground between evidentialists, who think that one's evidence determines what can be properly taken for granted, and non-evidentialists, who deny this. (On certain coherentist pictures, for example, it will make sense to talk about someone's evidence, but the fact that they have some evidence will be ultimately explained by patterns of coherence among their other beliefs, and will not be analytically prior to facts about rationality.) It would be convenient for several purposes if we could have an interest-invariant notion of evidence that explained why interests caused people to sometimes lose knowledge. Unfortunately, that's not a viable position. As I'll argue in Chapter 9, the arguments that knowledge is interest-relative generalise into arguments that evidence itself can be interest-relative.

So far, I've sketched in the very broadest outlines the kind of theory I'm going to propose. Here's the plan for how that theory will be laid out, and defended, over the coming chapters, as well as some more details on how the chapters relate to previously published work.

In Chapter 2, I'll set out the main argument for the interest-relative theory. The argument turns on how to think about a particular low-stakes bet. I argue that every option other than the interest-relative theory says very implausible things about this case.

The next two chapters set out the fundamentals of the theory. In Chapter 3 I lay out the interest-relative theory of belief, and how that view differs from the view I developed in "Can We Do without Pragmatic Encroachment?". Then in Chapter 4 I extend that to a theory of knowledge, and introduce a problem that will come up more in later chapters—how this theory interacts with closure principles.

The following three chapters are, in one way or another, responses to various objections to interest-relative theories. They are also the most novel parts of the book; very little of these three chapters draws on previously published work. Indeed, some of the key arguments build on work that was unpublished at least when I started work on this book. Chapter 4 draws heavily on work by Woodard (2020) and Chapter 7 draws heavily the doctoral dissertation of Nilanjan Das (2016).

In Chapter 5 I discuss the role that the concept of inquiry plays in my theory. In my theory, if something is known, it is available to use as

a starting point in inquiry. I used to think this meant I was committed to agreeing with Friedman (2019b) that it is incoherent to inquire into something one knows. I've come to see that this isn't right; depending on what one wants to do in an inquiry one may want to deliberately set aside some premises. That might mean inquiry into what one already knows is reasonable. This fact is used to respond to an influential objection by Jessica Brown (2008) to the style of argument I use in Chapter 2.

In Chapter 6 I respond to an objection that theories like mine are committed to implausible closure failures in cases where choosers have very similar options to choose between. There is a proof in "Can We Do without Pragmatic Encroachment?" that the theory developed there is immune to closure failures in these types of cases, so the objection can't be right as stated. It turns out that the reason the theory of that paper respects closure is that it has absurdly sceptical consequences in cases where there are similar objects to choose between. That's hardly better than a closure failure. In this chapter I aim to do better.

I show that the objection relies on the assumption that the chooser aims to maximise expected utility, and this isn't the right criteria of correctness for decisions in close call situations. It isn't true that when one is selecting cans off the supermarket shelf, one's selection is rational if and only if (henceforth, iff) it is utility maximising. Rather, the rational chooser in such a situation will adopt a strategy that has the best long-run consequences. In this case, the strategy will probably be something like the strategy of picking arbitrarily unless it is clear that one of the choices is defective. Given a theory of rational choice that emphasises the importance of decision-making strategies, rather than the importance of utility maximisation, my preferred epistemological theory gets the right answers. There are two traps to avoid here: closure failure and scepticism. And the focus on strategies lets us avoid both.

In Chapter 7 I respond to the frequently voiced objection that interest-relative theories lead to implausible verdicts about pairs of situations where knowledge is lost or gained due to what looks like an irrelevant feature of a situation. I have two responses to these objections. One was first offered in "Defending Interest-Relative Invariantism" (Weatherson, 2011). I argue that the intuitions are about what makes it the case that a person does or doesn't know something, and the argument from

these examples moves too quickly from a claim about modal variation to a claim about what knowledge is made of. The second response is, I think, more compelling, and it's essentially a point that Das makes which I borrow. These objections over-generate. Every modern theory of knowledge leads to pairs of cases where a person gains or loses knowledge depending on factors that seem 'irrelevant'. So it's not an objection to my view that it has the same consequences as every plausible theory of knowledge.

The last two long chapters go into relatively technical details of my theory of knowledge. I've put them at the end partially because they are technical—I don't want to lose readers until as late as possible! But also partially because they are the least changed from earlier work.

Chapter 8 goes over my theory of rational belief. Surprisingly, and in contrast to the view defended by Jeremy Fantl and Matthew McGrath (2009), interests affect rational belief in a very different way to how they affect knowledge. On my view, but not theirs, someone who has mistaken, and irrational, beliefs about what practical situation they are facing can easily have a rational, true belief that is not knowledge. This chapter also tidies up some loose ends from Chapter 3 concerning the so-called 'Lockean' theory of belief.

Chapter 9 sets out my interest-relative theory of evidence. I argue that one's evidence is what a Radical Interpreter would say one's evidence is. In some cases, this means we end up playing a kind of coordination game with The Radical Interpreter. What our evidence is turns on what the right solution to that game is. The solution is interest-relative, but not in the same way as knowledge, nor in the same way as rational belief.

Chapter 10 ends with a short note connecting interest-relativity to the familiar saying *Knowledge is Power*. I argue that this saying only makes sense given an interest-relative view of knowledge. If interest-relative theories were flawed for one reason or another, then we'd have to simply concede that the saying is false. We shouldn't concede that; the saying is true, and interest-relative epistemology explains why it is true.

