

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All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at 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

10. Power

Knowledge is power. That is, knowledge grounds our ability to do things. What things? Not, typically, bodily movements. If Lupin knows the passcode for the phone, he can unlock the phone. But even without that knowledge, he could make the bodily movement of typing in 220348 or whatever the passcode was. What power does he get from knowing the code is 220348? He gets the ability to deliberately unlock the phone. He gets the ability to unlock the phone by typing 220348, and for this to not be a lucky guess, but a rational action. Knowledge makes rational action possible. That's why it is powerful. That's what its power consists in. That's why the Nyāya philosophers were right to base anti-sceptical arguments on the possibility of rational action. Knowledge matters in everyday life; it explains why we have the power to act rationally.¹

But it turns out to be surprisingly hard to articulate the magnitude of that power. One might want to say that the actions which knowing that *p* makes rational include all the actions whatsoever that make sense if *p* can be taken as given. This, as we've seen many times over, can't be right. If it were right it would entail either that some actions that seem horribly reckless are in fact rational, or that an absurd form of scepticism is true, and almost no actions are in fact rational. If our pre-theoretic judgments of which actions are rational are even close to being right, there must be limits to the power of knowledge.

Are those limits sensitive to the interests of the would-be knower, or are they independent of those interests? I've argued that they are

1 The story I'm telling in this paragraph deliberately echoes the view that ability modals do not express possibility but necessity (Mandelkern, Schultheis and Boylan, 2017). To say Lupin can unlock the phone is not to say he might unlock the phone—anyone might unlock it hitting random numbers—it's to say he has a method that would unlock the phone if deployed.

sensitive to the knower's interests, against what I called the orthodox view that they are not. There are two primary reasons I've given for this.

One reason is that the interest-relative answer, unlike the orthodox answer, clarifies why the boundary between knowledge and non-knowledge matters. From the interest-relative viewpoint, the boundary between knowledge and non-knowledge is philosophically and practically significant. In contrast, the orthodox view treats the boundary between knowledge and non-knowledge as analogous to the distinction between heavy things and non-heavy things. How heavy something is matters a lot; but which side of the heavy/non-heavy boundary it falls on does not.

According to orthodoxy, to know something is to have enough power to use that knowledge to act rationally in, you know, a wide enough range of cases. How wide? It's sort of wide enough. Why is it this range rather than some other? Oh, no good reason. We just talk that way, and it's the job of epistemologists to explain why we do. No, not explain—because it's arbitrary and inexplicable. It's to describe the limits, limits which are ultimately quite arbitrary. This is no good at all. The heterodox interest-relative view has a better answer. The limits are set just where they need to be to make it true that what this person knows rationalises the actions that it should in fact rationalise.

A second reason for favouring the interest-relative view is that it makes some core principles of action theory and decision theory turn out to be correct, while the orthodox view makes them not quite right. Here are the (versions of the) principles I have in mind.

Means-End Rationality

If X should be aiming for end E , and X knows both that action A is the only means to end E , and that A will indeed lead to end E , to then X should intend to perform A .

Strict Dominance Reasoning

If X has to choose between A and B , and there is some partition P of a space of possibilities such that X knows both that precisely one member of P is actual, that X 's actions make no difference to which member is actual, and that conditional on each member of P , A is better than B , then X should prefer A to B .

Given orthodoxy, neither of these can be correct. In each case, X might know that A will lead to a good outcome, but it might also be probable

enough that *A* will lead to a disastrous outcome for it to be irrational for *X* to choose *A*. That's absurd, and so we should reject orthodoxy. The interest-relative view makes these, and any number of related principles, turn out unrestrictedly true.²

If these are the two reasons for adopting an interest-relative view, the appeal of some kind of contextualist version of an interest-invariant view should be clear. The contextualist can, and does, say that there is something special about the boundary between knowledge and non-knowledge. When one says that someone knows something, one means (more or less) that their evidence for that thing is good enough for one's own purposes. Any speaker of the language can truly say, "When I use 'knows', it means what I need it to mean, neither more nor less". We could quibble here, but that's something like a solution to the problem of arbitrariness.

The problems come with principles like Means-End Rationality and Strict Dominance Reasoning. These are both false within standard, interest-invariant, contextualist views.³ Contextualists are aware of this fact, and say that they have something nearly as good: a meta-linguistic replacement. Both principles are true if we replace talk about what *X* knows with talk about what *X* can truly self-predicate 'knowledge' of. It doesn't matter whether *X* knows that *A* will lead to end *E*, but whether *X* can say the words *I know that performing A will lead to end E*. But it's absurd to think that fundamental principles of rationality should make reference to a particular language, e.g., English, like this. So contextualism cannot get us out of this jam. This is not to say that contextualism is false. Maybe there are other reasons to believe in contextualism. It might be supported by general principles about the way that names behave in attitude ascriptions. Or it might be a consequence of the view that 'knows' is polysemous, a view I floated in a footnote in Section 2.7. Or it might be supported by reflection on cases where people talk about others with lower evidential needs than their own. I don't take a stance

-
- 2 Note that it's not much help to say that the principles are approximately true according to the orthodox picture. In game theory we want to be able to iterate principles like this, and approximately true principles can't be applied iteratively.
- 3 For the rest of this chapter, when I talk about contextualism, I mean the standard version of contextualism that does not entail that knowledge is interest-relative.

on these questions here; I'm just arguing that contextualism doesn't save orthodoxy.

All these benefits of the interest-relative view are of no use if the interest-relative view has even worse defects. And most of this book has involved pushing back against some of the alleged defects of the view. In some cases, I've avoided possible criticisms by adopting a form of interest-relative epistemology that doesn't allow the criticisms to take hold. Many existing critiques of interest-relative epistemology focus on forms of the view where being in a 'high-stakes' situation is relevant to what one knows. Since my version of the view does not say that interests matter iff they put the chooser in a high-stakes situation, those criticisms don't have any bite. Other critiques target a part-time version of interest-relative epistemology, where some but not all key notions are interest-relative. I used to endorse such a part-time interest-relative theory, but eventually decided that the criticisms of such a view were decisive.

Still, many criticisms do need replies, and much of this book has consisted in replying to them. In Chapter 3 I replied to (successful) criticisms of how I'd previously handled knowledge of propositions that are not relevant to the thinker's current inquiries. In Chapter 5 I replied to arguments that started with the observation that it sometimes seems right to double check things that we know. And in Chapter 6 I replied to arguments based on how some versions of interest-relative epistemology (including the version I'd previously endorsed) handled choices between nearly indistinguishable options.

Many criticisms of interest-relative epistemology do not turn on detailed engagement with how the view handles this or that case, but on the very idea of interests being relevant to epistemology. One way you see this is with arguments that just start from the implausibility of two people who are alike in evidence differing with respect to knowledge. But you also see it in the surprisingly common position in the literature that there is some extra-large burden of proof on defenders of interest-relativity. It seems to be a common presupposition that unless there is a super compelling argument that knowledge is interest-relative, we should reject the idea that it is, even if there is no equally compelling argument against the idea. I'm in general very sceptical of burden of proof arguments; it always looks like an attempt to win by bribing the umpires. But there is extra reason to be sceptical of these particular

burden of proof arguments. It's hard to even state a principle about the (alleged) independence of knowledge and interests without saying something false about double luck cases, or variable reliability cases, or deviant causal chain cases.

We've known since 1963 (if not many centuries before) that knowledge depends on more than just the evidential basis of the thing purportedly known. Since then there has been a dizzying array of proposals about what else knowledge might depend on. And on reflection, many of these proposals have been very plausible. This book aims to defend one such proposal, that knowledge depends on interests. In particular, how much evidential support one needs to have knowledge depends on what inquiries one is, and should be, engaged in. Given the tight relationships between knowledge and rational action, adding this to the vast list of things knowledge is sensitive to should not be surprising.

