

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All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at 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

Preface

Over the years I've written many papers defending an idiosyncratic version of interest-relative epistemology. This book collects and updates the views I've expressed over those papers.

My original plan was not a collection of papers, that would hardly add much value over a well-designed webpage, but a book that was largely structured out of different sections of different papers. My thought was that I had something like a working theory between the papers, and what would be useful would be to blend the sentences, paragraphs, and even whole sections from them into a coherent narrative. Some of that plan has been retained. Most sections in Chapters 8 and 9 are very similar to sections in one or other previously published paper. But the bulk of the book is new. In putting the pieces together, I realised that I'd changed my mind about enough things, and needed to express myself very differently about enough other things, so as to make it worth rewriting much of what I had. The result is that this is about 60% a new book, 20% a heavily edited version of previous material, and 20% lightly edited republishing of previous material. Even that last 20% has some value I think—it helps to see those points in the context of an overall story—but this is mostly a new book.

Interest-relative epistemologies all start in roughly the same way. A big part of what makes knowledge important is that it rationalises action. But for almost anything we purportedly know, there is some action that it wouldn't rationalise. I know what I had for breakfast, but I wouldn't take a bet at billion to one odds about it. Knowledge has practical limits. The first idiosyncratic feature of my version of interest-relative epistemology is how those limits are identified. Other interest-relative philosophers typically say that the limits have to do with stakes; in *high-stakes* situations knowledge goes away. That's no part of my view. I think knowledge goes away in *long-odds* situations. High-stakes

situations are almost always long-odds situations, for reasons to do with the declining marginal utility of money. But the converse isn't true. On my view, knowledge often goes away in cases where it is trivial to check before action. This idea, that interests matter in long-odds cases, and not just in high-stakes cases, is the main constant in what I've written on interest-relativity over the years.

But there are three other respects in which the view I'm going to set out and defend in this book is very different from the view I set out in older papers.

I used to say that knowledge went away in cases where replying on the purported knowledge would lead to the wrong answer. I focussed, that is, on the outputs of inquiry. Knowledge goes away if relying on it would lead one to make a mistake. I now think I was looking at the wrong end of inquiry. Knowledge goes away if the thinker starts conducting an inquiry where the purported knowledge is an inappropriate starting point, and inappropriate for the special reason that it might be false. One way to recognise a bad starting point is by realising that starting there will mean we end up at the wrong place. But it's not the only way. Sometimes a bad starting point will lead to the right conclusion for the wrong reasons. As the Nyāya philosophers argued, rational inquiry starts with knowledge. If it would be irrational to start this inquiry with a particular belief, that belief isn't knowledge.

Not all inquiries are practical inquiries, but many are. And practical inquiries are usually going to be at the centre of attention in this book. But what is someone trying to figure out when they conduct a practical inquiry? I used to think that they were trying to figure out which option maximised expected utility, and to a first approximation identified knowledge with those things one could conditionalise on without changing the option that maximised expected utility. As noted in the previous paragraph, I no longer think that we can identify knowledge with what doesn't change our verdicts. But more importantly, I no longer think that expected utility maximisation is as central to practical inquiry as I once did. There are theoretical reasons from game theory that raise doubts about expected utility maximisation being the full theory of rational choice. Weak dominance reasoning is part of our theory of rational choice, and can't be modelled as expected utility maximisation. Perhaps some kinds of equilibrium seeking are parts of practical

inquiry, and can't be modelled as expected utility maximisation. There are also very practical reasons to think that practical inquiry doesn't aim at expected utility maximisation. When there are a lot of very similar options—think about selecting a can from a supermarket shelf—and it's more trouble than it's worth to figure out which of them maximises expected utility, it's best to ignore the differences between them and just pick. As I'll argue in Chapter 6, this makes a big difference to how interests and knowledge interact.

In the version of interest-relativity that I'm defending here, everything in epistemology is interest-relative. Knowledge, rational belief, and evidence are all interest-relative. But they are all interest-relative in slightly different ways. The main aim here is to defend the interest-relativity of knowledge. A common objection to interest-relative theories of knowledge is that they can't be extended into theories of all the things we care about in epistemology. Here I try to meet that challenge. The way I do so is a little messy. It would be nice if some part of epistemology were interest-invariant, as I once thought, or if all the interest-relative notions were interest-relative in the same way, as other interest-relative epistemologists argue. For better or worse, that's not the view I'm defending. Interests matter throughout epistemology, and we just have to go case by case to figure out how and why they matter.

The ideas from the last three paragraphs are totally absent from my earliest work—several times they are explicitly rejected—but become more prevalent as the years go on. This is the first time I've defended them all in one place. And I think they are all necessary to make the theory I want to defend hang together.

Here is a list of these papers on interest-relativity that I've mentioned a few times already.

- Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19 (2005), 417–443.
- Defending interest-relative invariantism. *Logos and Episteme*, 2 (2011), 591–609.
- Games and the reason-knowledge principle. *The Reasoner*, 6 (2012), 6–7.

- Knowledge, bets and interests. In Jessica Brown and Mikkel Gerken (Eds.), *Knowledge Ascriptions* (75–103), Oxford University Press, 2012.
- Reply to Blackson. *Journal of Philosophical Research*, 46 (2016), 73–75.
- Games, beliefs and credences. *Philosophy and Phenomenological Research*, 92 (2016), 209–236.
- Reply to Eaton and Pickavance. *Philosophical Studies*, 173 (2016), 3231–3233.
- Interest-relative invariantism. In Jonathan Jenkins Ichikawa (Ed.), *Routledge Handbook of Epistemic Contextualism* (240–253), Routledge, 2017.
- Interests, evidence and games. *Episteme*, 15 (2018), 329–344.

I wrote most of this manuscript while on sabbatical at the Australian National University in the first half of 2019, and I’m very grateful for their hospitality while I was a visitor there.

Support for that sabbatical came from the Marshall M. Weinberg Professorship at the University of Michigan. And I’m once again incredibly grateful for the support Marshall has given to philosophy, and to many other disciplines, at the University of Michigan.

Many of the papers were drafted, and workshopped, while I was a Visiting Fellow at the Arché Research Centre at the University of St Andrews. You could probably fill a book this long with the mistakes I was talked out of in formal and informal meetings in St Andrews. And it was a real privilege to have been part of that community for a decade.

In Winter 2020 I taught a graduate seminar based off a draft of this manuscript at the University of Michigan. I received a lot of valuable feedback from the students in that seminar. I suspect I would have received even more valuable feedback had we not had to scramble to convert the course into a virtual event halfway through the semester. But I’m still very grateful for what I learned from them over that course.

I’ve presented this material at many departments and workshops, and am very grateful to the feedback I’ve received on each occasion. Most of the book was presented in one form or another at Arché. As well, parts have been presented at the 2012 Rutgers Epistemology Conference,

the 2017 *Episteme* Conference, a workshop on pragmatic encroachment organised by Arizona State University in 2017, the University of Sydney, the Australian National University, and the 2020 Ranch Metaphysics Workshop. I've also had valuable feedback on ideas in the book over the years from Michael Almeida, Charity Anderson, Thomas Blackson, Jessica Brown, Stewart Cohen, Josh Dever, Tom Donaldson, Tamar Szabó Gendler, Peter Gerdes, Katherine Hawley, John Hawthorne, Jonathan Ichikawa, Jon Kvanvig, Jennifer Lackey, Barry Lam, Harvey Lederman, Matthew McGrath, Sarah Moss, Jennifer Nagel, Shyam Nair, Daniel Nolan, Ángel Pinillos, Jacob Ross, Mark Schroeder, Kieran Setiya, Ernie Sosa, Levi Spectre, Robert Stalnaker, Jason Stanley, and Matthew Weiner. Jonathan Ichikawa also read over the whole manuscript and provided many useful comments.

And of course I've got more feedback, and more useful feedback, from Ishani Maitra than from anyone, or any place, else. She's had to listen to, and often talk me out of, any number of dead ends, false leads, and outright mistakes, on this topic for the best part of two decades. If there's anything in what follows that manages to be true, useful, and new, it's thanks to her feedback, advice, and support.

