Somayeh Tohidi
Curriculum Vitae

University of Oxford
Faculty of Philosophy
somayeh.tohidi@philosophy.ox.ac.uk

Areas of Specialisation

Formal Epistemology, Social Epistemology, Decision Theory, Moral Philosophy

Areas of Competence

General Philosophy of Science, Philosophy of Language, Feminist Philosophy, Islamic Philosophy

Employment

Research Associate, University of Oxford Oct 2026 – present
Foundations of Longtermism project, funded by the Leverhulme Trust and led by Prof. Richard Pettigrew

Research Associate, University of Bristol Feb 2026 – Sep 2026
Foundations of Longtermism project, until it moved to Oxford

Postdoctoral Researcher, University of Manchester Apr 2025 – Dec 2025

Academic Affiliations

Research Associate, London School of Economics, CPNSS Sep 2025 – present

Education

PhD Philosophy, London School of Economics 2020–2025
Passed without correction
Supervisors: Dr. Liam Kofi Bright; Prof. Richard Bradley
Examiners: Prof. Richard Pettigrew; Dr Jingyi Wu

Pre-doc Fellowship, Massachusetts Institute of Technology 2023–2024
Funded by SHASS
Supervisor: Dr. Kevin Dorst

MA Philosophy, Simon Fraser University 2017–2020
GPA: 4.06
Dissertation: *There is more to Probabilities than Meets the Eye: How Private Information can be Inferred from Announced Credences*
Supervisor: Dr. Nicolas Fillion

MSc Philosophy of Science, London School of Economics 2013–2015
Passed with distinction
Dissertation: *Verifiability Criterion and the Hard Limits of Science*
Supervisor: Prof. Bryan Roberts

MA Islam and the West: Comparative Studies, Middlesex University 2012–2013
Passed with distinction
Dissertation: *Mind and Body: A Comparative Study of Popper's and Mulla-Sadra's views*
Supervisor: Prof. Ali Paya

BSc Chemical Engineering, Sharif University of Technology 2005–2010

Teaching

Teaching Assistant

London School of Economics, London, UK

- Philosophy of Economics — Fall 2024, Winter 2025
- Philosophy and Society — Fall 2024
- Rationality and Choice — Fall 2022, Winter 2023
- The Big Questions: An Introduction to Philosophy — Fall 2021, Winter 2022

Simon Fraser University, Burnaby, Canada

- Critical Thinking — Fall 2017, Spring 2018
- Introduction to Logic and Reasoning — Spring 2017, Summer 2019

Facilitator

The Philosophy Foundation, London, UK

- Philosophy for Children: facilitating philosophical discussions in schools — 2015

Publications

1. **Unconvinced Yet Influenced: Vaccine Decisions in the Shadow of Misinformation** (with Ed Pertwee and Mihaela Popa-Wyatt) *(Forthcoming in Misinformation and Other Epistemic Pathologies, Cambridge University Press)*: This paper examines how misinformation shapes vaccine refusal, even while maintaining doubt towards false claims. Exposure to misinformation can alter one's perceptions of a vaccine's rare side effects and the perceived credibility of authoritative sources. When individuals are uncertain about vaccine safety, these changes can lead them to refuse vaccination. We use the expected utility framework to show that such refusal, while shaped by misinformation, can still be an instrumentally rational response under uncertainty rather than a cognitive failure. We identify three mechanisms through which misinformation influences decision-making: (1) erosion of trust in authoritative sources, (2) heightened emotional responses to anecdotal risks, and (3) a combination of relatively modest shifts in both trust and side-effect perception. Our model highlights the irreversible impact of misinformation and underscores the need for proactive interventions, such as pre-bunking and media literacy, to counter its effects.

Under Review

1. **Cognitive Economy and the Stability of Belief: The Case of Politically Salient Science** *(Erkenntnis; resubmitted after revise and resubmit)*: For a bounded Bayesian, working out their credence function over a more fine-grained partition is cognitively costly. On the other hand, such work can have epistemic benefits. In this paper, I first introduce a rational norm for how a bounded inquisitive Bayesian may decide whether to do that work, based on their information sources, that is to say, the information that they may receive in the *future*. Then I explore the consequence of this norm for the *categorical beliefs* that a rational agent may hold, given that some probabilistic accounts of belief, like the stability thesis, are partition-sensitive. I show that whether an agent should believe a proposition can depend on how likely they judge it that they will receive information settling whether a defeater of it is true.

   I apply this result to the case of a layman considering the claim of a scientific study. I argue that whether a layman should believe the claim of the study depends on how likely they think it is that they will receive information settling whether the study is biased. Their judgement of this probability in turn depends on whether the claim is politically salient and whether they consider themselves embedded in the political contestation around it. This conclusion has implications for disagreement over politically salient science: such disagreement *can be* rational, traceable to how agents assess their future information and not to their irrationality or to a difference in their current evidence.

In Preparation

1. **Is It Really Just About Scale? Catastrophe and the Longtermist Intuition**: Longtermists claim that preventing human extinction should take priority over preventing near-term harms, and they explain this by appealing to the vast number of future people affected. This paper challenges that explanation by designing a thought experiment, structured as an interventionist causal test, that contrasts the extinction case with the choice between revolution and reform under a corrupt government. I show that when the catastrophic nature of the outcome is removed while scale and ambiguity are held fixed, the clarity of the

longtermist preference weakens. This establishes that the catastrophic nature of extinction, not merely its scale, is a causal contributor to the longtermist intuition, with significant implications for the scope of longtermist recommendations.

2. **Rational Resistance to Correction: A Bayesian Analysis of the Continued Influence Effect** (with Ed Pertwee and Mihaela Popa-Wyatt): The continued influence effect (CIE) is a well-documented phenomenon whereby misinformation continues to affect people's beliefs even after they receive and accept a correction. Standard explanations attribute CIE to cognitive limitations or irrational tendencies. This paper demonstrates that CIE can occur even for ideally rational Bayesian agents, suggesting that the phenomenon is not merely a product of flawed reasoning but can arise from the structure of our informational environment. We model both misinformation and corrective information as testimony from sources that are not fully trusted. To formalize partial trust, we develop an account of defeater propositions within Bayesian epistemology. We distinguish two types of defeaters: inaccuracy (the source is careless or unreliable) and indeterminacy (the source is manipulative or deceptive). While uncertainty about inaccuracy defeaters permits standard conditionalization, uncertainty about indeterminacy defeaters does not: the rigidity condition fails, requiring an alternative updating method that we introduce and develop. Corrective information takes two forms: undercutting corrections target the source of misinformation (claiming it is inaccurate or insincere), while rebutting corrections directly contradict its content. We show that an ideal Bayesian can experience CIE after receiving either type of correction. This occurs even when the agent trusts the corrective source more than the source of misinformation: because neither source is trusted completely, both leave their mark on the agent's credences. We also show that timing matters for undercutting corrections about source manipulation: receiving such corrections before exposure to misinformation can prevent CIE in belief entirely, while corrections received afterward cannot.

3. **A Bounded Bayesian Has Rational Faith in You**: Decision making under uncertainty about members of demographic groups who are subject to negative stereotypes paves the way for a conflict between our moral and rational obligations. On one hand, we feel morally obliged to resist using demographic statistics about these groups in our decision making, as it can lead to behaviour that reinforces oppressive social practices against the group. On the other hand, we feel rationally obliged by the principle of total evidence to use all the evidence available to us, including demographic statistical evidence. In this paper, I argue that the conflict is avoidable for a bounded Bayesian, because a bounded agent, unlike an idealised one, has a genuine decision to make about whether to engage with a piece of evidence at all. I first offer a model for how a bounded Bayesian should make this decision: whether to engage with a piece of evidence and update their credence on it. I then use that model, together with Babic's model of updating on noisy statistical data, to show why such an agent has rational grounds to ignore stereotype-aligned demographic statistical evidence, and hence to avoid the conflict between morality and rationality. Finally, I show how the result extends to agents who are not statistically literate enough to appreciate the upshot of Babic's model.

4. **Choosing Moral Arguments in Social Movements: A Rational Choice Framework** (with Mihaela Popa-Wyatt): Social movements are fundamentally moral phenomena: they arise from perceived injustices and demand moral responses from potential participants. We argue that this moral character has strategic implications that standard models of social movements, which treat them as instances of the collective action problem, overlook. Research in moral psychology shows that individuals differ in their modes of moral reasoning. Because of these differences, the moral framing an activist chooses shapes how potential participants respond: those whose mode of moral reasoning matches the framing are more likely to be persuaded, while those whose reasoning mode is mismatched may remain unpersuaded or actively resist the message. We develop a rational choice framework that models potential participants as moral reasoners and analyzes the activist's strategic choice between deontological and consequentialist framing. The activist is uncertain about the distribution of moral reasoning types among potential participants. The framework generates two hypotheses: (1) deontological framing is more effective in early-stage movements, when emotional intensity is high and deep commitment is essential, while consequentialist framing is more effective in later-stage movements that require broader coalitions; (2) among consequentialist appeals, non-welfarist framing (emphasising justice and structural outcomes) is more effective than welfarist framing (emphasising well-being) at generating the expectation structures necessary to sustain participation. These hypotheses can be tested by future empirical research on movement communication strategies.

5. **Mind Your Probability Language**: This paper discusses the pragmatics of probabilistic statements about social groups. I use Grice's cooperative principle to argue that statistical reports about social groups whose content is aligned with a salient social stereotype implicate essentialist and causal claims

about those groups, and thereby contaminate the common ground of the conversation. I also show how, in certain policy-making conversational contexts, reporting a statistical correlation can implicate an intervention that reinforces oppressive social practices. Finally, I discuss how Gricean cancellation cannot mitigate these harms.

6. **I'm Not Even Uncertain About It. I Don't Think About It: Suspending Credence in Demographic Profiling**: Demographic profiling is usually discussed at the level of belief. Ross (2022) argues that profiling beliefs are intrinsically at odds with social equality, which is why an egalitarian should suspend judgment about whether a given person has some esteem-relevant trait, and why profiling is objectionable even when no action is at stake. I argue that the same worry arises one level down, at the level of credence, and that a bounded Bayesian has the resources to address it. For a bounded agent, forming a credence about a specific individual is not automatic. When the agent already holds a credence in a relevant proposition, such as a demographic base rate, forming a new credence about the individual requires fine-graining their partition and judging how the new proposition bears on what they already believe, while respecting probabilism. This is costly, and where it is costly the agent has a genuine liberty: they may decline to form the credence at all. I call this suspension of credence. I show when this liberty obtains, why demographic profiling is an instance of it, and that the relevant cases divide in two. Where no concrete decision about the individual is in prospect, social egalitarian values are enough to justify suspending credence about random members of a demographic group. Where a concrete decision is in prospect, egalitarian values alone are not enough, because they come into conflict with the instrumental demand to act well. For that second case I point to my paper in preparation, A Bounded Bayesian Has Rational Faith in You, which shows that in the most pressing instances, where a decision is in prospect and the statistics align with a negative stereotype, an agent can honour both demands without suspending credence, by declining to update the individual-level credence on the demographic evidence.

Invited Talks

1. "Rational Resistance to Correction", Munich Centre for Mathematical Philosophy, LMU Munich (July 2026)
2. "Mind Your Probability Language", Philosophy of Language Workshop, University of Edinburgh (June 2026)
3. "The Ethics of Foreign Intervention" (panel), London School of Economics (May 2026)
4. "Mind Your Probability Language", Eastern APA, Baltimore, MD (January 2026, cancelled)
5. "Unconvinced Yet Influenced: Vaccine Decision-Making in the Shadow of Misinformation", Trust & Trustworthiness Workshop, University of Leeds (May 2025)
6. "A Real Bayesian Has Faith in You", Northeastern Epistemology Workshop (NEW), Northeastern University, Boston (April 2025)

Selected Presentations

1. "Rational Resistance to Correction", BSTK, University of Oxford (September 2026)
2. "Choosing Moral Arguments in Social Movements: A Rational Choice Framework", University of Bristol (June 2026)
3. "Rational Resistance to Correction", Choice Group, London School of Economics (February 2026)
4. "Mind Your Probability Language", PPE, King's College London (July 2025)
5. "Mind Your Probability Language", Joint Session, University of Glasgow (July 2025)
6. "Probably Scientific, Definitely Political, I Don't Believe it", MIT-Michigan Workshop, MIT, Boston (April 2024)
7. "Stereotypes destabilize statistically-grounded credences", Formal Epistemology Workshop, Northeastern University, Boston (June 2023)

8. "Stereotypes destabilize statistically-grounded credences", Munich Centre for Mathematical Philosophy, Munich (March 2023)
9. "Generic Statements and Harmful Social Stereotypes", Athena in Action Workshop, Rutgers University, New Jersey (June 2022)
10. "Generic Statements and Harmful Social Stereotypes", Generic Generalizations in Science and Ethics, University of Oslo (June 2022)
11. "Stereotypes destabilize statistically-grounded credences", 17th Annual London-Berkeley Philosophy Conference, Senate House, London (May 2022)
12. "There is more to Probabilities than Meets the Eye: How Private Information can be Inferred from Announced Credences", International Conference for Philosophy of Science and Formal Methods in Philosophy, University of Gdansk (December 2019)
13. "There is more to Probabilities than Meets the Eye: How Private Information can be Inferred from Announced Credences", Canadian Society for the History and Philosophy of Science (July 2019)
14. "Putnam and the Reality of Time", 3rd LSE-Bayreuth Student Philosophy Conference (May 2015)

Awards and Distinctions

1. Predoctoral Fellowship, Department of Linguistics and Philosophy, MIT, 2023
2. Graduate Fellowship, Department of Philosophy, Logic, and Scientific Method, LSE, 2020
3. Graduate Fellowship, Faculty of Arts and Social Sciences, SFU, 2018
4. Graduate Entrance Scholarship, Faculty of Arts and Social Sciences, SFU, 2017
5. Ranked among top 1 percent in the National University Exam, Tehran, 2005

Service

1. Organizing a reading group with Sally Haslanger on 'How Change Happens' by Cass R. Sunstein
2. Co-organizer, LSE Choice Group seminar series
3. Judge, Ethics Cup (A programme held by University of St. Andrews for promoting rigorous discussion of ethics among high school students)

Technical Skills

- Programming: MATLAB and some Python
- Tools: LaTeX, Microsoft Office

References

Dr. Liam Kofi Bright, Department of Philosophy, Logic and Scientific Method; London School of Economics: liamkbright@gmail.com

Prof. Richard Bradley, Department of Philosophy, Logic and Scientific Method; London School of Economics: r.bradley@lse.ac.uk

Dr. Kevin Dorst, Department of Linguistics and Philosophy, Massachusetts Institute of Technology: kmdorst@mit.edu

Prof. Richard Pettigrew, Faculty of Philosophy, University of Oxford: richard.pettigrew@philosophy.ox.ac.uk