

Foundational Issues in Effective Altruism (October 2022 Version)

This syllabus is based on an 8-week course run at Oxford in Michaelmas 2017, led by [Hilary Greaves](#)—the original course can be found [here](#). It has been expanded to fill a 13-week semester, and some citations have been updated.

About this course

Effective altruism is a philosophy and social movement that uses evidence and reasoning to determine the most effective ways to benefit others. This course explores the philosophical foundations of this approach. We will see that this brings us into contact with a number of issues in philosophy, from epistemology, to the philosophy of mind, to the philosophy of science, and beyond.

Each week covers a single topic. To get a sense of how the topics hang together, we have grouped topics around a few core themes:

- Weeks 1-5: Who should we help, and how?
- Weeks 6-7: Individual Efficacy and Collective Action.
- Weeks 8-9: Epistemological issues facing Effective Altruists.
- Weeks 10-13: Should Effective Altruists embrace *longtermism*?

Target Audience

The course is intended undergraduates with an interest in ethics, policy, or decision theory. It is strongly advised that students have taken a prior course in ethics (or at least, enough prior courses in philosophy that they feel comfortable with a reasonably fast-paced introduction to some core concepts in ethics).

A useful book to introduce Effective Altruism and many of the issues that will be discussed throughout the course is William MacAskill's [Doing Good Better](#). It is recommended pre-reading.

How Should I Use This Material?

Short answer

You're welcome to use this material however you like.

Longer answer

We think that Effective Altruism is an important movement to study, and we think it raises substantive philosophical questions that students can engage well with. Our goal in putting sharing this syllabus is to help lecturers and students engage with the material, regardless of what form that ultimately takes. You might like to teach through the entire syllabus, incorporate individual topics into your own materials, or just use the reading lists for ideas and inspiration. It could also be used for self-guided study.

Giving feedback

We'd love to hear your feedback! This is a young research area, so we plan to revise this teaching material reasonably regularly (about once a year). We'd welcome all feedback on everything from the small (Did you spot a typo? Did we miss a key reading?) to the big (Did the topics work? Have we underrepresented a view?). You can submit feedback [here](#).

Topic 1: Duties in a World of Need

Motivating question: There is a great deal of suffering in the world, and some are much better off than others—how should we respond to those facts?

Core Readings

Peter Singer (1972), [Famine, Affluence, and Morality](#), *Philosophy & Public Affairs*, vol. 1, no. 3, pp. 229-243.

Thomas Pogge (2007), General Introduction and Section 9 of Chapter 4 in, [World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms, 2nd Edition](#), Polity, Cambridge.

Optional Readings

Peter Unger (1996), Chapter 2 in [Living High and Letting Die: Our Illusion of Innocence](#), Oxford University Press.

Onora O'Neill (2003), [A Kantian Approach to Famine Relief](#), in *Ethics: Contemporary Readings*, H Gensler, E Spurgin and J Swindal (eds.), Taylor and Francis.

Norbert Anwander (2005), [Contributing and Benefiting: Two Grounds for Duties to the Victims of Injustice](#), *Ethics & International Affairs*, vol. 19, no. 1, pp. 39-45.

Leif Wenar (2017), [Blood Oil: Tyrants, Violence and the Rules That Run the World](#), Oxford University Press. (Book-length, but a very good, accessible discussion of international resource privilege.)

Topic 2: Duties and Demandingness

Motivating question: Are there limits on how much we are required to sacrifice to aid those in need?

Core Readings

Kristen Bykvist (2010), [Is Utilitarianism Too Demanding?](#), in *Utilitarianism: A Guide for the Perplexed*, Continuum.

Travis Timmerman (2015), [Sometimes There Is Nothing Wrong With Letting A Child Drown](#), *Analysis*, vol. 75, no. 2, pp. 204-212.

Optional Readings

Liam Murphy (1993), [The Demands of Beneficence](#), *Philosophy & Public Affairs*, vol. 22, no. 4, pp. 267-292.

Brian McElwee (2015), [What Is Demandingness?](#), in *The Limits of Moral Obligation*, M. can Ackeren & M. Kühler (eds.), Routledge, New York.

William MacAskill, Andreas Mogensen & Toby Ord (2018), [Giving Isn't Demanding](#), in *The Ethics of Giving* P. Woodruff (ed.), Oxford University Press.

Benjamin Sachs (2019), [Demanding the Demanding](#), in *Effective Altruism: Philosophical Issues*, H. Greaves & T. Pummer (eds.), Oxford University Press.

Christan Barry & Holly Lawford-Smith (2019), [On Satisfying Duties to Assist](#), in *Effective Altruism: Philosophical Issues*, H. Greaves & T. Pummer (eds.), Oxford University Press.

Topic 3: Duties and Distance

Motivating question: Do we have the same duties to those far away from us as we do those in close proximity to us? Are we permitted to give special consideration to our nearest and dearest?

Core Readings

Frank Jackson (1991), [Decision-Theoretic Consequentialism and the Nearest and Dearest Objection](#), *Ethics*, vol. 101, no. 3, pp. 461-482.

Jeremy Waldron (2003), [Who Is My Neighbor?](#), *The Monist*, vol. 86, no. 3, pp. 333-354.

Optional Readings

Philip Pettit & Robert Goodin (1986), [The Possibility of Special Duties](#), *Canadian Journal of Philosophy*, vol. 16, no. 4, pp. 651-676.

Frances Kamm (2000), [Does Distance Matter Morally to the Duty to Rescue?](#), *Law and Philosophy*, vol. 19, no. 6, pp. 655-681.

Any article in *The Monist's* 2003 [special issue on moral distance](#).

Benjamin Lange (2022), [The Ethics of Partiality](#), *Philosophy Compass*, vol. 17, no. 8, pp. 1-15.

Topic 4: Who Matters?

Motivating question: If our goal is to do the most good possible, whose wellbeing should we focus on: humans or animals?

Core Readings

Peter Singer (2011), [Equality for Animals?](#), in *Practical Ethics (3rd Edition)*, Cambridge University Press.

Alastair Norcross (2004), [Puppies, Pigs, and People: Eating Meat and Marginal Cases](#), *Philosophical Perspectives*, vol. 18, no. 1, pp. 229-245.

Optional Readings

Jonathan Anomaly (2015), [What's Wrong With Factory Farming?](#), *Public Health Ethics*, vol. 8, no. 3, pp. 246-254.

Shelly Kagan (2018), [For Hierarchy in Animal Ethics](#), *Journal of Practical Ethics* vol. 6, no. 1, pp. 1-18.

Walter Sinnott-Armstrong & Vincent Conitzer (2021), [How Much Moral Status Could Artificial Ever Achieve?](#), in *Rethinking Moral Status*, S. Clarke et al. (eds.), Oxford University Press.

Ruth Faden, Tom Beauchamp, Debra Matthews & Alan Regenberg (2021), [Toward a Theory of Moral Status Inclusive of Nonhuman Animals: Pig Brains in a Vat, Cows versus Chickens, and Human-Nonhuman Chimeras](#), in *Rethinking Moral Status*, S. Clarke et al. (eds.), Oxford University Press.¹

¹ Interested students could read any other chapter in in *Rethinking Moral Status*.

Topic 5: Empathy and Identified Lives

Motivating question: Should we treat identified and statistical lives equivalently? Is it justified to prioritize identified lives over statistical lives on the grounds of empathy?

Core Readings

Dan Brock (2015), [Identified versus Statistical Lives: An Interdisciplinary Perspective](#), in *Identified versus Statistical Lives: An Interdisciplinary Perspective*, G. Cohen et al. (eds.), Oxford University Press.

Nir Eyal (2015), [Concentrated Risk, the Coventry Blitz, Chamberlain's Cancer](#), in *Identified versus Statistical Lives: An Interdisciplinary Perspective*, G. Cohen et al. (eds.), Oxford University Press.

Michael Slote (2015), [Why Not Empathy?](#), in *Identified versus Statistical Lives: An Interdisciplinary Perspective*, G. Cohen et al. (eds.), Oxford University Press.

Optional Readings

Any chapter in *Identified versus Statistical Lives: An Interdisciplinary Perspective*.

Larry Temkin (2022), Chapter 3 (especially 3.4-3.6) in [Being Good in a World of Need](#), Oxford University Press.

Topic 6: Making a Difference

Motivating Question: For many morally significant phenomena (e.g., factory farming or climate change), it seems like our individual choices do not make a difference. Is that true, and what (if anything) does that tell us about how we should contribute to addressing those phenomena?

Core Readings

Walter Sinnott-Armstrong (2005), [It's Not My Fault: Global Warming and Individual Moral Obligations](#), in *Perspectives on Climate Change: Science, Economics, Politics, Ethics*, Walter Sinnott-Armstrong and Richard Howarth (eds.), Elsevier.

Shelly Kagan (2011), [Do I Make a Difference?](#), *Philosophy & Public Affairs*, vol. 39, no. 2, pp. 105-141.

Optional Readings

Felix Pinkert (2015), [What if I Cannot Make a Difference \(and Know It\)?](#), *Ethics*, vol. 125, no. 4, pp. 971-998.

Julia Nefsky (2019), [Collective Harm and the Inefficacy Problem](#), *Philosophy Compass*, vol. 14, no. 4.

Mark Budolfson (2019), [The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response](#), *Philosophical Studies*, vol. 176, no. 7, pp. 1711-1724.

Julia Nefsky (2021), [Climate Change and Individual Obligations: A Dilemma for the Expected Utility Approach, and the Need for an Imperfect View](#), in *Philosophy and Climate Change*, M. Budolfson, T. McPherson & D. Plunkett (eds.), Oxford University Press.²

² Interested students could read any other paper in Section 3 of *Philosophy and Climate Change*.

Topic 7: Individual Efficacy and What We Do Together

Motivating question: Does Effective Altruism focus too much on individual difference-making? Should our focus instead be on institutions, states, or social structures?

Core Readings

Stephanie Collins (2019), [Beyond Individualism](#), in *Effective Altruism: Philosophical Issues*, Hilary Greaves and Theron Pummer (eds.), Oxford University Press.

Brian Berkey (2017), [The Institutional Critique of Effective Altruism](#), *Utilitas*, vol. 30, no. 2, pp. 143-171.

Optional Readings

Norbert Anwander (2005), [Contributing and Benefiting: Two Grounds for Duties to the Victims of Injustice](#), *Ethics & International Affairs*, vol. 19, no. 1, pp. 39-45.

Kok-Chor Tan (2010), [Rights, Harms, and Institutions](#), in *Thomas Pogge and His Critics*, A. Jaggar (ed.), Polity, Cambridge.³

Amia Srinivasan (2015), [Stop the Robot Apocalypse](#), *London Review of Books*, vol. 37, no. 18.

Alexander Dietz (2016), [What We Together Ought to Do](#), *Ethics*, vol. 126, no. 4, pp. 955-982.

Stephanie Collins & Holly Lawford-Smith, [The Transfer of Duties: from Individuals to States and Back Again](#), in *The Epistemic Life of Groups: Essays in the Epistemology of Collectives*, M. Brady and M. Fricker (eds.), Oxford University Press.

³ This picks up themes from a core reading from Topic 1: Thomas Pogge (2007), General Introduction and Section 9 of Chapter 4 in *World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms, Second Edition*, Polity, Cambridge.

Topic 8: Evidence, Interventions, and (Disputed) Efficacy

Motivating question: How can we know what kinds of interventions are effective, and how does this affect our prioritization among causes?

Core Readings

Toby Ord (2013), [The Moral Imperative Toward Cost-Effectiveness in Global Health](#), Center for Global Development (e-resource).

Julian Reiss (2013), 'Evidence-Based Policy' in [Philosophy of Economics: A Contemporary Introduction](#), Routledge.

Optional Readings

Dambisa Moyo (2008), [Dead Aid: Why Aid is Not Working and How There Is Another Way for Africa](#), Allen Lane, London. (The whole book is a readable introduction to aid-scepticism—Chapter 3 is a good starting point to get to grips with the core issues.)

Abhijit Banerjee and Esther Duflo (2011), [Low-Hanging Fruit for Better \(Global\) Health?](#) In *Poor Economics*, Penguin.

May Ann Bates & Rachel Glennester (2017), [Striking a Balance Between Theory and Action](#), *Sanford Social Innovation Review*, vol. 15, no. 3, pp. 48-54.

Interested students might like to look at [GiveWell's material on Randomized Controlled Trials](#). GiveWell is a large charity evaluator aligned with the Effective Altruist movement. (You can start at the link given and follow the links to explore further, depending on time and interest.)

Donal Khosrowi & Julian Reiss, [Evidence-Based Policy: The Tension Between the Epistemic and the Normative](#), *Critical Review*, vol. 31, no. 2, pp. 179-197.

Nancy Cartwright (2010), [What Are Randomised Controlled Trials Good For?](#), *Philosophical Studies*, vol. 147, no. 1, pp. 59-70.

Topic 9: Cluelessness

Motivating question: If we care about efficacy, then should we care about *all* of our acts' effects? Can we know what the total effects of our acts are, and if not, is that a problem?

Core Readings

James Lenman, [Consequentialism and Cluelessness](#), *Philosophy & Public Affairs*, vol. 29, no. 4, pp. 342-370.

Hilary Greaves, [Cluelessness](#), *Proceedings of the Aristotelian Society*, vol. 116, no. 3, pp. 311-339.⁴

Optional Readings

Elinor Mason (2004), [Consequentialism and the Principle of Indifference](#), *Utilitas*, vo. 16, no. 3, pp. 316-321.

Adam Elga (2010), [Subjective Probabilities Should be Sharp](#), *Philosophers' Imprint*, vol. 10, no. 5, pp. 1-11.

Susanna Rinard (2015), [A Decision Theory for Imprecise Probabilities](#), *Philosophers' Imprint*, vol. 15, no. 7, pp. 1-16.

J. Robert G. Williams (2016), [Indeterminacy, Angst and Conflicting Values](#), *Ratio*, vol. 29, no. 4, pp. 412-433.

⁴ Students unfamiliar with Expected Utility (or decision theory more generally) might want to read Ray Briggs' [Stanford Encyclopedia article on Expected Utility Theory](#). The Greaves reading in particular introduces *imprecise probabilities*: see Brian Weatherson's [Imprecise Decision Theory](#) or Section 4 of Lara Buchak's [Stanford Encyclopedia entry on Rivals to Expected Utility Theory](#).

Topic 10: Longtermism

Motivating question : The future potentially contains a vast number of people. If our goal is to do the most good, then should we shift our focus away from present generations and towards the (very) long-term future?

Core Readings

[Longtermism: The Future Is Vast - What Does This Mean for Our Own Life?](#), *Our World in Data*.

Hilary Greaves and William MacAskill (2021), [The Case for Strong Longtermism](#), *Global Priorities Institute Working Papers*, No. 5-2021.

Optional Readings

Nick Beckstead (2019), [A Brief Argument for the Overwhelming Importance of Shaping the Far Future](#), in *Effective Altruism: Philosophical Issues*, H. Greaves & T. Pummer (eds.), Oxford University Press.

Chapters 1-2 of [The Long View: Essays on Policy, Philanthropy, and the Long-Term Future](#) (Natalie Cargill and Tyler M John. Eds). FIRST Strategic Insight, London.

John Broome (1994), [Discounting the Future](#), *Philosophy and Public Affairs*, vol. 23, no. 2, pp. 128-156.

Simon Caney (2014), [Climate Change, Intergenerational Equity, and the Social Discount Rate](#), *Politics, Philosophy & Economics*, vol. 13, no. 4, pp. 320-340.

Topic 11: Population Ethics

Motivation question: How should we think about the value not just of individual lives, but populations as a whole? How much does population size matter, and what do we owe people who may or may not come to exist in the future? Do such questions affect how we should prioritise among causes?

Core Readings

Hilary Greaves (2017), [Population Axiology](#), *Philosophy Compass*, vol. 12, no. 11.

Nick Bostrom (2003), [Astronomical Waste](#), *Utilitas*, vol. 15, no. 3, pp. 308-314.

Optional Readings

Derek Parfit (1984), Chapters 16-19 in [Reasons and Persons](#), Clarendon Press, Oxford.

Elizabeth Harman (2004), [Can We Harm and Benefit in Creating?](#), *Philosophical Perspectives*, vol. 18, no. 1, pp. 89-113.

Jeff McMahan (2013), [Causing People to Exist and Saving People's Lives](#), *The Journal of Ethics* vol. 17, no. 1/2, pp. 5-35.

Gustaf Arrhenius and Torbjörn Tännsjö (2017), [The Repugnant Conclusion](#), in *The Stanford Encyclopedia of Philosophy*

Topic 12: The Long-term and Humanity's Trajectory

Motivating question: How should we think about the future of humanity—is it good or bad, getting better or worse?

Core Readings

William Macaskill (2022), Chapters 3 and 4 in [What We Owe The Future](#), *Basic Books*.

Optional Readings

Hanno Sauer, Charlie Blunden, Cicile Eriksen and Paul Rehren (2021), [Moral Progress: Recent Developments](#), *Philosophy Compass*, vol. 16, no. 19, pp. 1-10.

Elizabeth Anderson (2014), [Social Movements, Experiments in Living, and Moral Progress: Case Studies from Britain's Abolition of Slavery](#), *The Lindley Lecture for 2014*, University of Kansas, Department of Philosophy.⁵

Allen Buchanan & Rachell Powell (2018), [The Evolution of Moral Progress](#), Oxford University Press.⁶

William Macaskill (2022), Chapter 9 in [What We Owe The Future](#), *Basic Books*.

⁵ This reading discusses the treatment of slaves in a way that some people might find difficult. This can be largely (though not entirely) skipped by beginning at page 3.

⁶ Students could start with the Introduction, then move on to Chapters 1 and 4 for especially relevant discussion, then other chapters depending on time and interest.

Topic 13: The Long-term and Humanity's Survival

Motivating question: If our goal is to do the most good, should our top priority be minimising extinction risk? How does that goal compare to, say, improving the lives of people who currently exist?

Core Readings

Toby Ord (2020), Chapters 2 and 6 in [The Precipice](#), Bloomsbury.

Core Readings

Jason Matheny (2007), [Reducing the Risk of Human Extinction](#), *Risk Analysis*, vol. 27, no. 5, pp. 1335-1344.

Any other chapter in Ord's [The Precipice](#).

Roger Crisp (2021), [Would Extinction Be So Bad?](#), *New Statesman*.

David Thorstad (2022), [Existential Risk Pessimism and the Time of Perils](#), *GPI Working Papers Series*, no. 1-2022.