

Michaelmas term 2017: Foundational issues in effective altruism

Professor Hilary Greaves

Logistics

This course is offered under the rubric of Philosophy Paper 198, “Special subjects”. It is open to FHS students in all joint schools including Philosophy. Paper 103 (“Ethics”) is strongly recommended as a prerequisite for this course, as is William MacAskill’s book “Doing Good Better” (OUP 2015).

Teaching

This course will be taught via lectures and classes. Attendance at both is mandatory. See below for times.

You *must* also read the ‘core readings’ for each topic before the corresponding class, and you may *prefer* to read them before the lecture. I will collect as many of possible of the readings in electronic form, and will upload them to a dedicated Weblearn site for the course (to be placed in Philosophy → Teaching and learning material → FHS → Michaelmas Term 2017).

The classes will be centred around student presentations. For each class, you should do both of the following:

- 1) **Group exercise:** Divide the students in your class into two groups. Each group must contain at least 2 students, and should (after group discussion) prepare a short (20 minute) presentation on a focussed topic related to the lectures and readings for the week in question. Over the course of the term, each student should present the equivalent of at least two full class presentations.

Each class will have a student coordinator for each class, to settle the division of classes into groups for presentation. For the first week’s class, the coordinators are as follows; please decide among yourselves who will be your class’s coordinator after that.

Class 1 (Fridays 10.30-12.20): Zachary Leather

Class 2 (Fridays 2.30-4.30): Natalie Beckett

The format of the presentation should be much like that of a tutorial essay: you should

- a. articulate a clearly defined question;
- b. in attempting to answer this question, carefully consider arguments on both sides; and
- c. defend a reasoned view.

The class venue will probably not have Powerpoint facilities (TBC), but you can prepare a paper-based handout and/or use a flipchart/whiteboard. You could also prepare Powerpoint slides to display on your own laptop (I’ll confirm later whether or not we will have access to a data projector, but until further notice, assume not).

I include some *suggestions* for questions that you might like to base a presentation alongside the reading list for each week, but you need not feel constrained by these suggestions if you have other ideas.

I also include some suggestions for how to give a good class presentation, as an appendix to this document.

- 2) **Individual exercise:** Make a note of questions that you would like to have clarified during class discussion. We'll start each class by pooling these questions, and will try to answer/discuss as many as possible during the course of the class.

There is no written work associated with the classes.

Lectures and classes

Lectures: These will normally be on Mondays, 4-5 pm. The exceptions are week 1 (Tuesday 4-5pm) and Week 5 (Monday 3-4pm).

Class 1 (10.30am-12.30pm): georgia.bacon@st-hildas, jack.baillie@wahdam, romil.depala@st-annes, zachary.leather@balliol

Class 2 (2.30-4.30pm): natalie.beckett@wadham, david.coates@exeter, aidan.goth@oriel, juliet.pursaill@new, emily.tench@bnc

Examination by extended essays

There will be no closed-book examination ("three-hour exam") for this paper. Instead, your grade for the course will be based exclusively on two extended essays. These should *total* 10,000 words, *unless* you are a MP/PP/CSP 4th year student, in which case they should total 15,000 words.

You can discuss possible essay topics briefly with me during the course of the classes this term. Subsequently, you will have two individual 30-minute tutorials in Hilary Term 2018, to provide written and verbal feedback on first drafts of each of your two extended essays. These tutorials will be with Ketan Ramakrishnan. You must submit your first essay draft by Friday of 0th week (HT18), and your second essay draft by Friday of 4th week (again HT18); please email them to Ketan directly (ketan.ramakrishnan@philosophy.ox.ac.uk). This teaching schedule applies regardless of whether you are offering this paper for examination in 2018 or instead in 2019.

Extended essays must be submitted for examination by Friday of Week 1 in the term of examination (for most of you, this is Trinity Term 2018; for a couple of you it is Trinity Term 2019). Essays must be handed in in hard copy at the Examination Schools, in an envelope marked "Essays for Philosophy special subject: Foundational issues in effective altruism". The essays must be clearly marked with your candidate number, and *not* your name or any other identifying information. They will be double-marked as part of the standard (double-blind) Oxford examination system.

Week 1: Moral obligations towards the global poor

Reading list

Core readings (mandatory)

Singer, Peter. "Famine, Affluence, and Morality." *Philosophy & Public Affairs* 1.3 (1972): 229-43 ([ejournal](#)).

Unger, Peter K. *Living High and Letting Die: Our Illusion of Innocence*. New York; Oxford: Oxford UP, 1996 ([ebook](#)).

[Chapter 2: "Living High and Letting Die: A Puzzle About Behavior Toward People in Great Need"]

Pogge, Thomas. *World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms*. Cambridge: Polity, 2002.

[Introductory chapter ("General introduction"), and section 9 of chapter 4 ("The causal role of global institutions in the persistence of severe poverty")]

Jaggar, Alison M (ed.). *Thomas Pogge and His Critics*. Cambridge: Polity, 2010.

[Chapter 2: "Rights, Harm, and Institution", sections 1 ("Pogge's institutional approach") and 4 ("Does the global order unjustly harm the poor when it disadvantages them?")]

Further readings (optional)

Anwander, Norbert. "Contributing and Benefiting: Two Grounds for Duties to the Victims of Injustice." *Ethics & International Affairs* 19.1 (2005): 39-45 ([ejournal](#)).

Entry on 'Libertarianism' in the (online) *Stanford Encyclopaedia of Philosophy* ([e-resource](#)).

Cruft, Rowan. "Human Rights and Positive Duties." *Ethics & International Affairs* 19.1 (2005): 29-37 ([ejournal](#)).

Wenar, Leif. *Blood Oil: Tyrants, Violence, and the Rules That Run the World*. New York: Oxford UP, 2017 (and alternative edition). (ebook ordered)

[Book-length, but a very good, accessible, extended discussion of the 'international resource privilege']

Suggested questions for class assignments

- Is there any morally relevant difference between failing to donate to aid the global poor and failing to rescue a child from a shallow pond?
- Can any reasonable moral theory escape the conclusion that affluent Westerners ought to donate significant amounts of money to wherever will do the most good?
- Can a duty to help the global poor be derived from the negative duty not to harm?
- Suppose that my government inflicts unjust harm on the citizens of a poor country. Do I bear partial responsibility for this harm because the government is my agent, because I benefit from the harm, because I contribute to the harm, for some other reason, or not at all? (Particularly relevant here: the article by Anwander on the 'further readings' list.)

Week 2: The importance of effectiveness, and evidence for effectiveness

Reading list

Core readings (mandatory)

Ord, Toby. *The Moral Imperative toward Cost-Effectiveness in Global Health* (2013): Policy File ([e-resource](#)).

Freedman, David, Robert Pisani, and Roger Purves. *Statistics*. 4th Ed. New York; London: W.W. Norton, 2007 (and alternative editions).

[Chapters 1 and 2]

Bates, Mary Ann, and Rachel Glennerster. "Striking a Balance Between Theory and Action" *Stanford Social Innovation Review* Summer 2017: 48-54 ([ejournal](#)).

The discussions of RCTs on GiveWell's blog: start [here](#), and follow the links to probe deeper (depending on time and interest)

Further readings (optional)

Easterly, William. *The White Man's Burden: Why the West's Efforts to Aid the Rest Have Done so Much Ill and so Little Good*. Oxford: Oxford UP, 2006 [on aid scepticism].

Moyo, Dambisa. *Dead Aid: Why Aid Is Not Working and How There Is Another Way for Africa*. London: Allen Lane, 2008 (legal deposit [ebook](#)) [another classic on aid scepticism].
Variant title: *Dead aid : why aid makes things worse and how there is another way for Africa*.

Shadish, William R., Thomas D. Cook, and Donald T. Campbell. *Experimental and Quasi-experimental Designs for Generalized Causal Inference*. Belmont, California, 2002, pp. 83-93 [on external validity].

West, Stephen, Jeremy Biesanz and Steven Pitts, "Causal inference and generalization in field settings: Experimental and quasi-experimental designs", in Harry T. Reis and Charles M. Judd (eds.). *Handbook of Research Methods in Social and Personality Psychology*. Cambridge: Cambridge UP, 2000 [another overview of experimental designs, complementary to Freedman et al]. ([ebook](#))

Cartwright, Nancy. "Understanding and Misunderstanding Randomized Controlled Trials." National Bureau of Economic Research (2016) ([e-resource](#)) [a relatively advanced discussion urging against overestimating the merits of RCTs].

Claridge, Jeffrey, and A. Fabian. "History and Development of Evidence-based Medicine." *World Journal of Surgery* 29.5 (2005): 547-53 ([ejournal](#)) [some background on the medical case – as the title suggests!].

Solesbury, William. "Evidence based policy: Whence it came and where it's going", working paper, available [here](#) [a similar recent-historical account of the rise of evidence in public policy].

Suggested questions for class assignments

Suppose that some highly cost-effective intervention is available to boost health and school attendance, but (i) only works on boys, (ii) only works on members of the ethnic majority, or (iii) is only feasible in urban areas. To what extent do equity considerations mitigate the cost-effectiveness-based case for prioritising this intervention over one that is less cost-effective but applicable to a broader group of beneficiaries?

Does the greater availability of RCTs justify prioritising global poverty interventions over (e.g.) interventions aimed at reducing existential risk?

"The insistence on backing up recommendations with RCTs is appropriate for a charity evaluator, whose task is specifically to issue recommendations to third parties, but the donors themselves should not penalise causes for which RCTs aren't available." Discuss.

Given the problem of external validity (= "the generalizability puzzle"), is there any advantage to randomisation in the development context?

Week 3: Making a difference

Core reading

Parfit, Derek. *Reasons and Persons*. Oxford: Clarendon Press, 1984 ([ebook](#)).
[Chapter 3: "Five Mistakes in Moral Mathematics"]

Kagan, Shelly. "Do I Make a Difference?" *Philosophy & Public Affairs* 39.2 (2011): 105-41 ([ejournal](#)).

Sinnott-Armstrong, Walter. "It's not my fault: Global warming and individual moral obligations", in Sinnott-Armstrong, Walter, and Richard B. Howarth. *Perspectives on Climate Change: Science, Economics, Politics, Ethics*. Amsterdam; San Diego, CA: Elsevier JAI, 2005, pp. 285-307 ([ebook](#)).

Broome, John. *Climate Matters: Ethics in a Warming World*. New York: W.W. Norton, 2012, pp. 74-78.

Optional further reading

Parfit, 'What we together do'

Nefsky, Julia. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan." *Philosophy & Public Affairs* 39.4 (2011): 364-95 ([ejournal](#)).

Arntzenius, Frank, and David McCarthy. "Self Torture And Group Beneficence." *Erkenntnis* 47.1 (1997): 129-44 ([ejournal](#)).

Pinkert, Felix. "What If I Cannot Make a Difference (and Know It)." *Ethics* 125.4 (2015): 971-98 ([ejournal](#)).

Dietz, Alexander. "What We Together Ought to Do." *Ethics* 126.4 (2016): 955-982 ([ejournal](#)).

Stephanie Collins (2013) Collectives' Duties and Collectivization Duties, *Australasian Journal of Philosophy*, 91:2, 231-248 ([ejournal](#))

Collins, Stephanie. & Lawford-Smith, Holly. "The Transfer of Duties: from Individuals to States and Back Again", in Michael Brady and Miranda Fricker (eds.). *The Epistemic Life of Groups: Essays in the Epistemology of Collectives*. Oxford: Oxford UP, 2016, pp. 150-172 ([ebook](#)).

Budolfson, "Collective Action, Climate Change, and the Ethical Significance of Futility". Available from the author's [website](#)

Suggested questions for class assignments

- Are there any cases in which a sufficiently large group can make a difference, yet no individual can? If so, what is the significance of such cases for ethical theory?
- In what sense, if any, does the problem of climate change require collective rather than individual action?

Week 4: Aggregative and non-aggregative approaches to value

Core readings (mandatory)

Taurek, John M M. "Should the Numbers Count?" *Philosophy & Public Affairs* 6.4 (1977): 293-316 ([ejournal](#)).

Otsuka, Michael. "Saving Lives, Moral Theory, and the Claims of Individuals." *Philosophy & Public Affairs* 34.2 (2006): 109-35 ([ejournal](#)).

Voorhoeve, Alex. "How Should We Aggregate Competing Claims?" *Ethics* 125.1 (2014): 64-87 ([ejournal](#)).

Further readings (optional)

Scanlon, Thomas. *What We Owe to Each Other*. Cambridge, Mass.; London: Belknap Press of Harvard UP, 1998.

[Chapter 5 ("the structure of contractualism"), part 9 ("aggregation"); pp. 229-241 in the paperback edition]

Kamm, F. "Aggregation and Two Moral Methods." *Utilitas* 17.1 (2005): 1-23 ([ejournal](#)).

Halstead, John. "The Numbers Always Count." *Ethics* 126.3 (2016): 789-802 ([ejournal](#)).

Norcross, Alastair. "Comparing Harms: Headaches and Human Lives." *Philosophy & Public Affairs* 26.2 (1997): 135-67 ([ejournal](#)).

Carlson, Erik. "Aggregating Harms - Should We Kill to Avoid Headaches?" *Theoria* 66.3 (2000): 246-55 ([ejournal](#)).

Suggested questions for class assignments

- Why, if at all, should we save the many rather than the one in Taurek's puzzle case?
- Is Scanlon right to insist that one must interrupt the TV broadcast to rescue the technician, regardless of the number of viewers?
- How much of a problem is it for the "Aggregate Relevant Claims" view (described by Voorhoeve) that that view sometimes recommends choosing the foreseeably worse outcome?
- Deworming treatments provide a *modest* benefit to *very large numbers* of individuals. For the same cost, distribution of insecticide-treated bed nets would provide a large benefit to a much smaller number of individuals. Does this consideration count in favour of bed net distribution over deworming, relative to the evaluation that would be suggested by straightforward utilitarian aggregation?

Week 5: Identified vs. statistical lives

Core readings (mandatory)

Brock, Dan. "Identified versus Statistical Lives: An Interdisciplinary Perspective", in Glenn Cohen et al. (eds.), *Identified versus Statistical Lives: An Interdisciplinary Perspective*. New York: Oxford UP, 2015, pp. 43-50 ([ebook](#)).

Eyal, Nir. "Concentrated risk, the Coventry blitz, Chamberlain's cancer", in Glenn Cohen et al. (eds.), *Identified versus Statistical Lives: An Interdisciplinary Perspective*. New York: Oxford UP, 2015, pp. 90-110 ([ebook](#)).

Verweij, Marcel. "How (not) to argue for the rule of rescue", in Glenn Cohen et al. (eds.), *Identified versus Statistical Lives: An Interdisciplinary Perspective*. New York: Oxford UP, 2015, pp. 137-149 ([ebook](#)).

Slote, Michael. "Why not empathy?", in Glenn Cohen et al. (eds.), *Identified versus Statistical Lives: An Interdisciplinary Perspective*. New York: Oxford UP, 2015, pp. 150-156 ([ebook](#)).

All of these are in Glenn Cohen et al. (eds.), *Identified versus Statistical Lives: An Interdisciplinary Perspective*. New York: Oxford UP, 2015 ([ebook](#)).

Further readings (optional)

The remainder of the essays in the Cohen et al volume.

Suggested questions for class assignments

- Why, if at all, should one prioritise identified over statistical lives? (The maximally broad question for this week's topic.)
- "Empathy pushes us towards prioritising identified over statistical lives. But empathy is known to track merely superficial features, and is thus morally irrelevant; we should resist its guidance." Discuss.
- Can an anti-aggregationist approach justify prioritising identified over statistical lives?
- Should consideration of identified vs. statistical lives have induced Churchill to take defensive action against the Coventry raid, in the scenario Eyal discusses?
- Consider a variant of Singer's 'shallow pond' example, in which one could save more people by selling one's suit (and then donating the resulting money to maximally cost-effective global poverty charities) than by jumping into the pond (ruining one's suit in the process). Can any reasonable moral theory hold that one ought not to save the drowning child from the pond in such a case? What are the implications of your answer for the debate over identified vs. statistical lives?

Week 6: Varieties of altruism

Core readings (mandatory)

Batson, Daniel et al. "Empathy and altruism", in C. R. Snyder, and Shane J. Lopez (eds.). *Oxford Handbook of Positive Psychology*. 2nd ed. Oxford: Oxford UP, 2009, pp. 417-427 ([ebook](#)).

Elster, John. "The Valmont effect: The warm-glow theory of philanthropy", in Patricia M. L. Illingworth, Thomas Pogge, and Leif Wenar (eds.). *Giving Well: The Ethics of Philanthropy*. New York; Oxford: Oxford UP, 2011, pp. 67-80 ([ebook](#)).

Snowden, 'Donating to multiple charities'

Macaskill, William. "Replaceability, Career Choice, and Making a Difference." *Ethical Theory and Moral Practice* 17.2 (2014): 269-83 ([ejournal](#)).

Further readings (optional)

Batson, C. Daniel. *Altruism in Humans*. New York: Oxford UP, 2011 ([ebook](#)).

[Esp. chapter 5 and appendices B-G (on the empirical evidence concerning the “empathy-altruism hypothesis”)]

Oakley, Barbara A. *Pathological Altruism*. Oxford; New York: Oxford UP, 2012 ([ebook](#)).

Suggested questions for class assignments

- Is it rational to split one’s philanthropic donations between more than one charity? Is such behaviour consistent with true altruism?
- “Working in the charity sector indicates a somewhat altruistic personality, but earning to give manifests a much purer form of altruism.” Is that so?
- Must someone who behaves altruistically for the sake of the resulting ‘warm glow’ be deceiving herself? Is the same true of all impure forms of altruism?

Week 7: Cluelessness

Core readings (mandatory)

Lenman, James. "Consequentialism and Cluelessness." *Philosophy & Public Affairs* 29.4 (2000): 342-70 ([ejournal](#)).

Greaves, Hilary. "XIV—Cluelessness." *Proceedings of the Aristotelian Society* 116.3 (2016): 311-39 ([ejournal](#)).

White, Roger. “Evidential symmetry and mushy credence”, in John Hawthorne and Tamar Gendler (eds.). *Oxford Studies in Epistemology. Volume 3*. Oxford: Clarendon, 2005. Oxford Studies in Epistemology, pp. 161-188. (soon to be available from WebLearn)

Further readings (optional):

Rinard, Susanna. "A Decision Theory for Imprecise Probabilities." *Philosophers' Imprint* 15.7: 1-16 ([ejournal](#))

Elga, Adam. "Subjective probabilities should be sharp." *Philosophers' Imprint* 10.5: 1-11 ([ejournal](#)).

Williams, J. Robert G. "Indeterminacy, Angst and Conflicting Values." *Ratio* 29.4 (2016): 412-33 ([ejournal](#)).

Suggested questions for class assignments

- To what extent, if at all, does the type of cluelessness focussed on by Lenman (and called ‘simple cluelessness’ by Greaves) handicap consequentialist decision-making? Is ‘complex cluelessness’ different, and if so how and why?
- Is it rational to refrain from donating on the basis of worries about cluelessness?
- What truth is there, if any, in the principle of indifference?

Week 8: Population ethics and existential risk

Core readings:

Greaves, Hilary. ‘Population axiology’, forthcoming in *Philosophy Compass*. ([Preprint](#))

Beckstead, ‘The case for existential risk reduction’. (Available from WebLearn.)

Bostrom, N. "Pascal's mugging." ([Link](#))

Optional further reading:

Parfit, Derek. *Reasons and Persons: Part IV*. Oxford: Clarendon Press, 1984 ([ebook](#)).
[Chapters 16-19 inclusive]

For the literature in population ethics more generally: follow up the references in the above article by Greaves.

Suggested questions for class assignments

- To what extent should the Repugnant Conclusion lead us to reject a totalist population axiology?
- Outline the argument for prioritising existential risk reduction. Is this argument sound?
- Should we accept expected utility theory even in cases that involve *extremely* low probabilities of *extremely* high stakes?

Appendix: Some tips for giving good class presentations

Structure. You should have a clear map of the logical structure of your presentation. For example, it might naturally fall into four sections, where section 1 clarifies the question and sets aside some potential misunderstandings, section 2 proposes an argument, section 3 considers possible objections to that argument, and section 4 briefly sums up. (This is only one possible structure, of course; in general, you should identify the appropriate structure after working out what you want to say, rather than expect your message to fit any particular standard template – every presentation is different!) Having worked out how to organise your material into a clear structure, *communicate that structure clearly to your audience* – this massively helps the audience to understand what you are saying. For example, it is generally a good idea to start the presentation by laying out the structure, and then explicitly announce when you are moving on to the next section.

What do you really need to say, and what don't you need to say? You are bound to have numerous thoughts on the topic, including interesting ones well worthy of discussion, that are (however) relatively tangential to the central thread of your presentation. It is generally better to *exclude* these: your presentation will be much clearer if you stick to the *central* points. This also helps with time-keeping: if you try to include every thought you have, you will almost certainly struggle to keep your presentation within the allotted time limit. (If you really want to flag more peripheral issues for possible discussion, you can include pointers to them as a list of very brief bullet points in a postscript to your main presentation, if time permits.)

Timekeeping. Most presenters run over time. Whatever the context of your presentation (from class assignment to job interview, etc.), *it is important not to do this*. Practice your presentation in advance *and time it*, so that you know how long it will take you to cover the ground you plan to cover, and you have an opportunity to re-organise or cut material as necessary. This also gives you a chance to 'hear out loud' which parts of your presentation don't sound as clear as you'd like, and re-think those parts. When actually giving the presentation, have some form of stopwatch with you, so that you can check your progress (by now you should know how long each sub-section of your talk

ought to take). (FWIW: When preparing professional talks for academic conferences, I usually find (i) that my first version is around 50-75% over the time limit, and (ii) that after cutting the talk down to meet the time limit my talk is much clearer, more streamlined, and generally better.)

Teamwork. You might decide to split the presentation between two or more members of your group. This is fine, but if you do, make sure the division of labour is clear both to the presenters and to the audience (e.g. one person taking on each 'section' of the presentation in turn), and make sure each person knows their personal time limit (i.e. you need to have co-ordinated so as not to collectively overrun the overall time limit).

Visual aids. Consider what level and type of visual aid (if any) might best help your audience to follow what you are saying, while also being compatible with your personal delivery style. Possibilities include pre-printed handouts and real-time annotation on a whiteboard/flip-chart, and, if the venue has AV facilities, Powerpoint presentations; in the case of pre-prepared visual aids, you have a choice of levels of detail (from brief bullet points to fuller summary). Try to avoid including so much detail on the visual aid that you end up simply reading from it – it is generally better to keep the visual aid relatively simple, and then talk 'around' it.

Delivery. Make eye contact with your audience, and make sure you are not speaking (i) too fast or (ii) too quietly for the audience easily to make out what you are saying. Try not to read your presentation from a pre-prepared script – this tends to make the delivery 'wooden'. Try to avoid verbal ('umm') and physical (playing with your hair) tics. Try to communicate (in your tone and body language) that you care about the topic.

Stage fright. If you find presenting to an audience nerve-wracking, you are not alone. Try to stay calm, and if you're not calm, try not to show that. Yes, this is all a lot easier said than done. One trick I used to use was reading out (verbatim from a script) the first paragraph or so of my presentation, because I found that after getting started, I was able to relax into it a bit more. If you feel you've made a mistake during your presentation, don't dwell on it over-much: it's fine (and helpful) to briefly correct it, but remember that the 'mistake' will seem much less significant to the audience than it perhaps feels to you. Other tips I've seen include paying attention to your breathing, and pausing for a sip of water now and again. The good news is that this gets a *lot* better with practice.

There is also plenty of advice on all this online- e.g. try running a search on 'tips for giving good presentations'.