

# Contemporary Moral Theories: Introduction to Contemporary Consequentialism Syllabus [Kian Mintz-Woo]

## Course Information

In this course, I will introduce you to some of the ways that consequentialist thinking leads to surprising or counterintuitive moral conclusions. This means that I will focus on issues which are controversial such as vegetarianism, charity, climate change and health-care. I have tried to select readings that are both contemporary and engaging. Hopefully, these readings will make you think a little differently or give you new ideas to consider or discuss.

This course will take place in SR09.53 from 15.15-16.45. I have placed all of the readings for the term (although there might be slight changes), along with questions in green to help guide your reading. Some weeks have a recommended reading as well, which I think will help you to see how people have responded to the kinds of arguments in the required text. Please look over the descriptions and see which readings you are most interested in presenting for class for our first session on 10.03.2016. (Please tell me if you ever find a link that does not work or have any problems accessing content on this Moodle page.) I have a policy that cellphones are off in class; we will have a 5-10 minute break in the course and you can check messages only during the break.

As a bonus, I have included at least one fun thing each week which relates to the theme of that week.

## Marking

The course has three marked components:

1. a weekly **presentation** of the main text by one or two students (20%),
2. **weekly short responses** (200-300 word) to the main or recommended texts (due Tuesday at noon before class) (all together 20%),
3. and a **final paper** (60%).

If you are presenting that week, you **do not need to have a short response!** There will be no penalty if you miss up to a *combination* of two things between weekly responses and classes missed. For instance, you can miss two classes or you can miss one weekly response and one class. However, if anyone attends **all** the classes **and all** of the weekly responses, I will round them to the nearest grade up. Also, I realize that people have different levels of background and I will try to keep this in mind while marking.

For **presentations**, either a handout (1 page either single- or double-sided) or a slideshow (presentation/'beamer') is required. If you want to use a slideshow, you can email it to me (at least three days before class at noon) or you can bring it on a USB drive, but please tell me either way. If you want to bring a handout, please print enough copies for the group. Your presentation does not and should not completely explain everything in the paper. Part of the goal of the presentation is to be able to select some part(s) of the paper that you think are interesting (and I'll be pleased if you also manage to use the recommended readings in your analysis, but this is not required).

Every presentation should

1. choose 2-3 ideas from the text that you think are interesting or important;
2. explain why those are interesting or important--how do these ideas support the author's view? Do these ideas change how you think about something? How do you think these ideas should influence society?;
3. explain why you agree or disagree--is the argument valid (are the links between the different parts of the argument logical)? Are there counterexamples to the claims? Could opposing positions be raised?; and
4. include at least three questions about the text to help start the discussion.
5. The presentation cannot be longer than 15 minutes, but can be shorter--the goal is just to start the class discussion, not to be complete!

The **short responses** have three functions. First, they are to make sure that you get credit for reading the texts before coming the class. Second, they are to help you learn to read like a philosopher. When a philosopher reads a text, she usually is thinking about which of the claims she agrees with and which she objects to and which she is interested in. In the short response, I want you to do these things: pick one or two ideas from either the main text or a recommended text and explain what they are, why they are interesting/surprising/(un)convincing/new to you, and why you (dis)agree with them. I prefer responses which are critical and show something that you think is mistaken in the text, but it is of course okay to also explain why you agree with the idea(s). You do not need to copy the text, you can cite it by page number. For instance, you could write "Bostrom argues that we should maximize the probability of an 'okay' outcome and calls this the maxipok rule (19). I think that Bostrom's defence of the maxipok rule is mistaken because..." Finally, if there are particularly interesting questions or thoughts in short responses, I may use them to advance the discussion in class (although I will not say who they are from unless the author wants me to).

These can be submitted in the assignment sections on Moodle at noon (12:00) on Tuesday before each Thursday class. The first one is due at noon on 15.03 for our class on 17.03. They can be submitted late up to the day of class, but if submitted late, I may not be able to return comments before class and your mark for that assignment will be lower than otherwise.

The hope is that one of your short responses (or your presentation!) can lay the groundwork to think about a topic for a final paper.

The **final paper** will be discussed later in the course, but I want you to be thinking about potential topics that interest you. I am happy to meet with anyone who has an idea of an argument they want to make either for or against any of the papers in class (including the recommended readings) after the course ends. I will discuss this further on during the course. Some resources for writing a philosophy essay can be read under the final class at the bottom of this Moodle page.

### Introduction (10.03.2016)

No prereading required, but if you listen to the podcast, that would be very helpful.

- BBC Podcast on Utilitarianism (2015): <http://www.bbc.co.uk/programmes/b05xhwqf>
  - BBC's Radio 4 has a podcast on Utilitarianism--it's a great introduction. You can download the MP3 version to listen on the go or can listen on the BBC webpage. It

features one of my old teachers, Brad Hooker. Janet Radcliffe Richards is quite amusing in it as well.

- Carla Bagnoli on How to Write a Philosophy Paper:  
<https://pantherfile.uwm.edu/cbagnoli/www/paperguidelines.html>
  - This is an excellent short introduction on writing philosophy papers.
- XKCD – Being Wrong: <https://xkcd.com/386/>
- \*Optional: A Pretty Silly (Non-Academic!) Introduction to Utilitarianism:  
<http://raikoth.net/consequentialism.html#utilitarianism>

Does consequentialism make demands on us for the global poor? (17.03.206 / Philipp Schönegger)

After discussing the Singer text, we will have a Giving Game, where we discuss three potential charities to donate money to... and we will actually donate money (ten euros/class attendee) according to the majority of how the class decides after discussion.

To make for a reasonable discussion, there are three charities that I have chosen. These are [Innovations for Poverty Action](#), [Oxfam](#), and [Against Malaria Foundation](#).

- Peter Singer – *Practical Ethics* Chapter 8 (and 1)
  - Peter Singer is the most influential utilitarian working in ethics today, although he has changed the type of utilitarian he is throughout his career (he began as a hedonist, meaning that he thought utility was pleasure/happiness, and now he is a preference utilitarian, meaning he thinks utility is the satisfaction of desires or preferences). This is his book on applied ethics. We will not be reading the whole thing, although if you find it interesting, it is available free through this link as an e-book at our library. (If the link below does not work, you can just find it in the Suchportal UB Graz as Practical Ethics.)

The required reading is Chapter 8, although I also recommend reading Chapter 1, especially if you have not done much ethics before. In Chapter 1, Singer discusses what ethics are in a university setting. In Chapter 7, Singer discusses a revised version of a very famous argument that he developed in a paper called *Famine, Affluence, and Morality* (you can download that article here: <http://www.jstor.org/stable/2265052>). In both the original version and the revised version in Chapter 8, Singer argues that, in light of the huge amount of suffering from poverty in the world, rich people (which include most of the people living in the developed world, and certainly most of the philosophy audience that read ethics papers) are allowing people to die and that, if allowing someone to die is not intrinsically different from killing, we are all murderers. He argues that this means we should donate significant amounts of money to proven charities which decrease poverty and suffering.

Many people find this conclusion very strong, and you may too if you have not read this before. What are your personal reactions? Do you think that philosophical arguments like this can motivate people? Should they motivate people? Why? Think through the steps and see if you agree with all of the ways the argument is developed. If you don't agree, explain where you think he has gone wrong. If you do agree, think about why other people wouldn't or what new information he has given you.

- Recommended – Murphy – “The Demands of Beneficence” (1993)
  - In this recommended reading, Liam Murphy argues against Singer's claims as over-demanding (this is often called the "Over-Demandingness Objection"), since many people do not in fact come close to complying with their duties of beneficence to the poor. Murphy suggests (a) that we could limit our duties by our "agent-centered prerogatives"--meaning our personal projects and relationships, or (b) that our duties of beneficence are limited to those which are compatible with ideal moral conditions under which everyone else fulfills their duties as well. Since our duties if everyone else also fulfilled their own duties would be much smaller (e.g. we would have to donate far less to the poorest if everyone else donated the same), Murphy claims that our duties are lower than Singer suggests.

How do you think Singer would respond? Does it matter that our current situation is non-ideal? Should we establish moral principles by appealing to the actual world or the ideal world? Why? How important would our own personal projects have to be to outweigh the demands of the extremely poor?

- Recommended – Charlotte Newey – “Self Serving Biases” (2016, §4-6)
  - Newey argues in a forthcoming paper (very contemporary!) that we should be wary of arguments that show we have limited duties to the global poor because we may be rationalizing our (in)actions and be subject to self-serving biases. This is a very helpful check in this debate (as well as many others) and it is important for philosophers to consider whether the arguments they advance help their own interests, not least because they should be more critical of such arguments. I recommend reading the most relevant sections: sections 4-6.

Do you notice arguments in daily life where different participants argue about which is fair, but different determinations seem to help the participants proposing those determinations? Do you notice such arguments in philosophy? How do you think one could determine whether one is motivated by a self-serving bias? Does it apply in the case of Singer's aid requirements?

- Assignment: Singer (Poverty)
- AntonyWH/Singer - Tweets (2016): <https://twitter.com/AntonyWH1/status/701879783593005056>
- Rolli - Selfishness (2015): <https://twitter.com/rolliwrites/status/663129807878668288>
- Existential Comics - A Dialogue on Morality: <http://existentialcomics.com/comic/91>
- [Link to] Certificate for our classes' Giving Game donation!

### What does Utilitarianism Presume? (14.04.2016 / Eleanor Young)

- Krister Bykvist – “Utilitarianism in the 20th Century” (2014)
  - Before we begin, it's worth discussing some of the modern background to utilitarianism (which is a version of consequentialism--although the predominant version), because the term is liable to be misunderstood. Bykvist (who was one of my teachers) is very clear that there are two main theses that make up utilitarianism: Act Consequentialism and Sum-Ranking Welfarism. Understanding these is very important to being able to follow consequentialist developments. Note that, if one just accepts AC and not SRW one is still a consequentialist, but not a utilitarian (in the chapter, Bykvist does talk

about what other forms of consequentialism would result).

[Do not worry if you cannot follow all of this chapter; it is more technical than the other texts in the course. The reason I include it is because if you want to understand precisely what is meant by utilitarianism or consequentialism, I think it is good to have a formal definition to appeal to when you get stuck. And the chapter might make more sense when you look back at it after the course is done.]

What is the difference between direct and indirect utilitarianism? Does welfare seem like the most plausible consequence (how does it differ from well-being)? What is the type of well-being that you think is most important? Why? Does maximization seem like the best pattern of consequentialism? Bykvist writes "Welfarism thus excludes information about desert, merit, freedom, justice, and rights (assuming that well-being is not even partly constituted by these factors)" (107). What does this mean? Do you think well-being is ('even partially') constituted by these factors? Is the axiomatic method helpful to you? Why or why not? How does utilitarianism differ from prioritarianism? Try to see if you can follow Harsanyi's argument--but don't worry if you find it challenging. Do you find Rawls' or Barry's objections to Harsanyi's argument convincing? Why or why not? If you have questions about some part of the chapter, you can describe them, their importance, and why you have them in the short response.

- \*Optional: Roger Crisp –“ Equality, Priority and Compassion”
- Assignment: Bykvist
- SMBC – Utilitarian Society: <http://www.smbc-comics.com/index.php?id=3857>
- SMBC – Introspection: <http://www.smbc-comics.com/index.php?id=3501>

### Should Consequentialists use Cost-Effectiveness Analysis in Healthcare? (Wolfgang Dorfinger)

- Peter Singer – “Why We Must Ration Health Care” (2009)
  - This argument is targeted towards Americans, but the point applies generally. Singer argues that we must ration health care because we cannot afford to treat everything without limit. He suggests a common consequentialist method, cost-benefit analysis. This method relies upon determining the probabilities and costs of various health outcomes. To many people, this appears morally problematic. One of those people (Frances Kamm) is the writer of the following paper.

Do you think Singer is right that we must ration health care? If so, do you think cost-benefit analysis is the right way to do it? Why or why not? If not, where do you think Singer goes wrong (it might be worth reading Frances Kamm's opposing argument to see some possible ways)?

- Recommended – Frances Kamm – “Cost Effectiveness Analysis and Fairness” (2015)
  - In this recommended reading, Kamm worries about aggregation in the ethics of health in response to claims like Singer's. This is an important topic in modern consequentialism (and anti-consequentialist critiques). Remember that there are different forms of consequentialism. Utilitarians count utility as mattering morally (see Krister Bykvist's chapter), and then aggregation is simple: one sums over utility. Some

who worry about consequentialism, like Kamm, worry that this addition is too simple: one intuition is that utility matters depending on who gets it. Some think that it matters more to those worse off--they are called prioritarians (these are also consequentialists!). But it is important to remember that utility is not some stuff that is handed around--it is not a benefit itself; it is a way of measuring benefits.

What is the difference between 'cost-effectiveness' and 'cost-benefit analysis'? (If you don't know, you can search for these terms and learn.) Do you think it matters who gets some utility (or 'benefit')? [Remember that, if we use the word 'benefit' or 'outcome' instead of 'utility', 'benefits' or 'outcomes' do not mean actual goods that can be distributed, they mean the positive or negative value to people of those goods (e.g. we do not mean food, but we mean the pleasure or satisfaction that people can get from consuming food).] What do you think fairness means? Do you think that someone who is more ill should get priority in terms of utility over someone who is less? Why or why not? Do you think the same about money (that a person who is more ill should get more money spent on them)? What is the difference between the previous two questions? Is it possible that there is some number of minor harms (e.g. toothaches) such that enough of them matter the same amount as a great harm (e.g. cancer)? Why or why not?

- Recommended – Arrow et al. – “Is there a Place for Cost-Benefit Analysis in Environmental, Health, and Safety Regulation?” (1996)
  - In this short (but dense) *Science* piece, economists Arrow et al. suggest various conditions under which cost-benefit analysis should be used in public policy.

Do you think that cost-benefit analyses do help decision makers? Under which circumstances would they help and under which would they hinder? Suppose that cost-benefit analyses are (i) absent or (ii) present in some department (e.g. environment)--which do you think would be better and why? Can you see problems with (i) or with (ii)? Their eighth condition suggests that distributional concerns matter (i.e. where or to whom benefits and costs go). Do you agree and why? Does this do enough to satisfy the concerns of Kamm (previous recommended reading)?

- Assignment: Singer (Healthcare Cost-Benefit)
- Existential Comics – Utility Monster: <http://existentialcomics.com/comic/8>
  - A utility monster is someone who derives more utility from any object you can give her or him than anyone else. Thus, utilitarianism says you should give everything to the utility monster (as long as they stay a utility monster regardless of what you give them).

## Should Animals Count as Much as Humans? (28.04.2016 / Ursula Byrne & Romana Stjepanovic)

- Peter Singer – Practical Ethics (Ch 3)
  - In this chapter, Singer discusses animals in the context of utilitarianism and introduces his famous term 'speciesism', which is meant to suggest that elevating humans above non-human animals is a mere prejudice like the prejudices some have for certain human 'races' (racism) or sexes/genders (sexism). Singer's argument is that it is capacity to suffer (sometimes called 'sentience') that is morally relevant, and that suffering matters regardless of the sufferer. Most of the time--we believe--humans can suffer more than non-human animals, but Singer claims in some cases this may not be true. More

importantly, if Singer's claim is true, then the non-human interests in not suffering mean we should treat animals very differently than we do now. For instance, intensive farming is responsible for more than 97% of the meat available in the United States. In Europe, there are more protections for animals (for instance they have the Five Freedoms), but many of the animal products available are not local so are not subject to EU regulation. We will not be reading the whole book, although if you find it interesting, it is available free through this link as an e-book at our library. (If the link below does not work, you can just find it in the Suchportal UB Graz as Practical Ethics.)

Is it reasonable to compare privileging species with privileging race or sex? Why or why not? Is this a mere prejudice or can it be defended by appeal to particular principles? Why does Singer think sentience is of such fundamental moral importance? Do you agree? Do you think most people agree when they eat meat or do you think they have reasons or responses? What do you think is the best response to Singer's position? Does this make utilitarianism too demanding?

- Recommended – Goodpaster – “On Being Morally Considerable” (1978)
  - In this text, Goodpaster asks whether sentience (or the ability to feel pain) is morally important. His suggestion is that sentience is just an evolutionary adaptation that certain living beings have, but that it should not be thought of as morally significant in light of this.

Does it matter that one is a moral person, or just potentially a moral person? Is Singer right that beyond sentience 'there is nothing to take into account' or is Goodpaster right that there is something much more fundamental—life? Can you explain Feinberg's argument that rights (or 'moral considerability') be grounded in moral persons? Why does Goodpaster disagree? Do interests presuppose desires or wants or aims? Why or why not?

- Recommended – Shelly Kagan – “What’s Wrong with Speciesism?” (2015)
  - In this very recent lecture (it was given last year), Shelly Kagan challenged Singer's claim that speciesism is a prejudice. Kagan argued that Singer does not establish that speciesism is a prejudice, unlike racism or sexism. In the latter cases, Kagan suggests, there are false factual beliefs (e.g. that races have different social or mental attributes), while the same is not true with speciesism. Kagan develops a view that fits his own intuitions better which he calls modal personism. On this view, beings which have the capacity or possibility ('modality') of becoming moral persons matter

What are the bases of racism and sexism? Are they factual or moral beliefs? Do you think that speciesism is a prejudice? How do you think Singer would respond to Kagan's objections (as a bonus, you can check by reading Singer's actual response!)? Do you believe in modal personism? Do you think that it captures the right intuitions? Can you think of cases which it does not capture?

- Recommended – Peter Singer – “Why Speciesism is Wrong: A Response to Kagan” (2015)
  - This short response to Kagan's piece above allows us to see how Singer regards Kagan's objections. His view is that there is an argument against speciesism, and this is based on the relationship with racism and sexism. (As you can see, he also refers to his text Practical Ethics, which we have looked at already.)

Do you think that Singer's response to Kagan is convincing? Do you think that Kagan

would be convinced? Why does it matter that 'there is something to be like' to morally matter? If there is 'nothing to be like' a plant or a car, does that mean that plants or cars do not matter? Does Singer give a convincing response to Kagan's modal personism theory? Do you think modal personism is more or less plausible after reading what Singer has to say about it?

- Assignment: Singer (Speciesism)
- SMBC – Are Humans Conscious?: <http://www.smbc-comics.com/?id=2867>
- SMBC – Center of the Universe: <http://www.smbc-comics.com/?id=2302>
- \*Optional – Sylvia Killingsworth – “Why Not Eat Octopus?” (2014): <http://www.newyorker.com/tech/elements/eating-octopus>
- \*Optional – David Foster Wallace – “Consider the Lobster” (2004): [http://www.gourmet.com/magazine/2000s/2004/08/consider\\_the\\_lobster.html](http://www.gourmet.com/magazine/2000s/2004/08/consider_the_lobster.html)
  - Wallace is perhaps the most brilliant and amusing writer I have ever read. His style is very unusual--lots of digressions and footnotes--but his thinking and vocabulary are incredible. In this essay, he goes to a festival for lobster in Maine (in Northeastern USA), and discusses the festival as well as the meaning of killing and eating lobsters. I highly recommend trying this piece when you have some time free.

## What are the Very Worst Consequences? (12.05.2016 / Daniel Pözlner & Bernhard Geißler)

- Nick Bostrom – “Astronomical Waste” (2003)
  - In this short paper, Bostrom argues that uncolonized space is a waste of resources. This is facilitated by some simple calculations on the possibilities for energy and resources that are not being consumed by space-faring humans or other sentient creatures.

Does this make utilitarianism look more plausible or less plausible? An opportunity cost is the cost of not pursuing the best course of action (e.g. the opportunity cost of not living in a house you own could be the most rent you could get from renting it to others). What kind of response do you have to the claim that our opportunity cost in not technologically advancing (or space colonizing) is so high? What are the utilitarian assumptions Bostrom makes (and are they needed)? Are you convinced that existential risk is the most important goal? Why or why not? Can the utilitarian really simplify the standard from maximize expected utility to minimize existential risk? If we adopt a person-affecting view (a view where what matters depends on the people that exist, ignoring potential people), how does Bostrom's argument change? Do you think 1% is a plausible estimate? Do you think that sufficient technological progress could defeat diminishing marginal returns for consumption goods?
- Nick Bostrom – “Existential Risk as a Global Priority” (2013)
  - Bostrom suggests that global policy (and individual efforts) should be towards mitigating existential risks, that is risks that threaten humanity as a whole. Consequentialism leads to conclusions like this because the consequences of human extinction (if we think that the relevant consequences are largely or wholly anthropocentric or human-focused) are about the worst thing that consequentialists can consider. However, notice that the question of what the right means are is not that important to Bostrom (such means are much more interesting to non-consequentialists, especially to deontologists). He is interested in preventing this very bad outcome and in figuring out how bad this outcome is relative to other very bad outcomes.

Some of the existential risks are things that you may not have considered before: for instance, he worries about dangerous artificial intelligence. He also suggests that 'aging' is a problem that can be solved. Do you think these are serious problems? Why or why not? Do you think that consequentialists should concern themselves with very unlikely problems, or just ones that seem to be more immediate (e.g. climate change)?

- Assignment: Bostrom (Existential Risks)
- SMBC – Robots: <https://www.smbc-comics.com/index.php?id=2124>
- SMBC – Altruism: <https://www.smbc-comics.com/index.php?id=2956>
- \*Optional – *New Yorker* – Bostrom Profile:  
[www.newyorker.com/magazine/2015/11/23/doomsday-invention-artificial-intelligence-nick-bostrom](http://www.newyorker.com/magazine/2015/11/23/doomsday-invention-artificial-intelligence-nick-bostrom)
  - This is a profile for the *New Yorker* about Nick Bostrom and his institute at Oxford FHI (Future of Humanity Institute). You can read this, entirely for fun, to hear about some of the weirdness of being one of the very influential consequentialists alive today. A couple people I know make a cameo appearance, including the excellent Niel Bowerman.

### How do Consequentialists Think About Climate Change? (19.05.2016 / Theresa Wakonig & Jonas Geike)

- Dale Jamieson – “When Utilitarians Should be Virtue Theorists” (2007)
  - Virtue theory tells us that the moral practice or way of life is to cultivate virtues. Most of the virtues lie in between some range of vice (too much anger in your character is a vice, too little is also a vice, so the virtue is the 'mean' between the two). However, this structure does not tell us very much about which virtues to cultivate or even what a virtue is, except that it is a trait or characteristic that can be cultivated in one's personality which tends to promote good (virtuous!) behaviour. Virtue theory is perhaps the oldest of the predominant moral systems (deontology/consequentialism/virtue ethics), since Aristotle endorsed it near the beginning of what we think of as the Western philosophical tradition (interestingly, similar views were endorsed in Eastern philosophical traditions as well, and we can find similarities to Aristotle in Confucian thought). However, different authors have interpreted this theory in light of different virtues that seem relevant in their time. Dale Jamieson continues in this tradition by introducing green virtues that he thinks should be cultivated by people to become more respectful and caring of the environment. This is important given that we are knowingly causing climate change, but many of us do not desire to make the problem worse. The relevant link to this course is that he thinks that cultivating these virtues can be justified on consequential grounds, in particular, on utilitarian grounds (remember, utilitarianism is a form of consequentialism!).

How does Jamieson get from utilitarianism to endorsing green virtues? Do you think this is convincing on utilitarian grounds? Does endorsing a different moral theory on utilitarianism weaken or strengthen utilitarianism? Is Jamieson right that 'non-contingency' is an important aspect of utilitarian responses to climate change? Is the utilitarian calculus as an action-guiding method as rejectable as Jamieson claims? How does progressive consequentialism compare with utilitarianism? Do you think that the green virtues Jamieson endorses are the right ones? The wrong ones? Why?

- Recommended – John Broome – “The Ethics of Climate Change” (2000)

- In this short popular (for the public) article, John Broome introduces the ethics of climate change and the role of economic thinking in determining how much we should spend to mitigate (reduce the impacts of) climate change. Although it is very short, it carefully lays out some of the important issues of climate change from a utilitarian perspective.

One important distinction that Broome introduces is between 'pure' discounting and consumption discounting. The first applies to any benefits or utility that occurs in the future (and has less importance) and the second applies only to social monetary (or 'consumption') benefits. Why is there this difference? Can you see reasons why one would be discounted and the other not? (Usually, consumption discounting is much larger than pure discounting and contains it, since consumption benefits are a kind of benefit which gets discounted in the 'pure' manner.) Does the market rate matter? Should it? What arguments can you see for referring to market behavior instead of reported preferences (or 'stated preferences')? What arguments can you see against it?

- Assignment: Jamieson
- Joel Pett – Better World: <https://www.gocomics.com/joelpett/2009/12/13/>
- \*Optional – Naomi Oreskes – “The Collapse of Western Civilization”
  - This is entirely optional, but I think it's a fascinating thing to read and suggest you at least try a page. In this short piece of futuristic fiction, Oreskes and Conway consider how the future might go if climate change is very bad and people decide to geoengineer. Trying to understand how we look from the future might make our actions on climate change look very different. It has a few jokes in it, too.

## Do Scientific Experiments Show that Consequentialism is True? (02.06.2016 / Antonia Webb)

- Joshua Greene – “Beyond Point-and-Shoot Morality” (2014)
  - One of the recent (past 10-15 years) trends in ethics and philosophy more generally is experimental philosophy (or x-phi). The idea is that, while philosophy seems to have depended on the intuitions of philosophers, perhaps it is important (or at least relevant) what and how non-philosophers think or intuit. Experimental philosophers tend to address this problem by doing a selection of the following: (a) examining the non-philosophical empirical literature, e.g. psychology or cognitive neuroscience, to see if these help philosophical inquiry; (b) proposing or implementing their own social scientific studies--sometimes with experimentally trained scientists--to test what non-philosophers think about philosophically interesting cases; or (c) arguing that experimental philosophy fills in a gap that traditional 'armchair' philosophy (without examination of empirical literature) fails to do.

Do you think Greene's article shows that neuroscience matters? Does Greene think that the neuroscience shows morally significant results alone? Do you agree and why or why not? Do you think this article shows that philosophers should be interested in these other fields? Why or why not? What does Greene think is the relevant place for empirical results? Do you think this works? What is the structure of his argument that consequentialist reasoning is more stable or should be favored over non-consequentialist reasoning? Is it convincing and why?

- Recommended – Selim Berker – The Normative Insignificance of Neuroscience (2009)

- This is not required reading, but is highly recommended as a critical discussion of Greene's approach and argument. Berker argues that the neuroscientific and empirical content is relatively (or perhaps completely) unimportant, and that it is the normative assumptions that drive Greene's argument.

What are the empirical claims Greene needs for his argument? What are the non-empirical (or normative) claims that Greene needs? How well does Berker's reconstruction of Greene match up with the original article? Do you think that Greene would be happier to take the 'shoddy inference' or the 'substantive normative intuitions' route (or do you think Greene would take a different interpretation altogether)?

- Assignment: Greene
- McSweeney's – "Lesser Known Trolley Variations": <http://www.mcsweeneys.net/articles/lesser-known-trolley-problem-variations>
  - Read McSweeney's Internet Tendency regularly. It is very very funny. Warning: swearing and coarse language.

#### How do Consequentialists Think About Careers? (09.06.2016 / Victoria Gastón Guiu)

- Will MacAskill – "Replaceability, Career Choice, and Making a Difference" (2013)
  - Deciding on a career is something that many of us do not spend a large amount of time thinking about. MacAskill points out that this is odd--we can have a substantial impact by considering what type of work we choose to do.

Have you thought about what career you want to have? Have you thought about what the impact of you going into this career could be (as opposed to another)? Does this approach make career choice too algorithmic? What are the values that should be introduced into a job search?

- Assignment: MacAskill
- SMBC – Effective Superman: <http://www.smbc-comics.com/?id=2305>
- New Yorker – Locally Grown: <https://condenaststore.com/featured/a-sign-for-produce-in-a-grocery-store-reads-roz-chast.html>

#### Objection 1: Does Consequentialism Lead to Paralysis? (16.06.2016 / Eunji Kang)

- James Lenman – "Consequentialism and Cluelessness" (2000)
  - Do we know the consequences of our actions? Not all of the consequences. We usually have a pretty good idea of the immediate consequences (although not always even these). Lenman argues that this should strike paralysis into consequentialists in this essay: some action might lead to very distant horrible (or great!) consequences which make the immediate or predictable consequences of the action completely unimportant in comparison. Standard utilitarianism says that the utility for any individual from any time matters as much as any other, so these distant consequences should matter as much as close ones (some consequentialist theories do not say so, such as 'discounted utilitarianism', which is popular in economics and is the subject of my own research). In the light of our ignorance of the far future impacts of our actions, how can consequentialism guide us? This is a modern version of a famous objection to consequentialism which is often called 'lack of action-guidingness' or 'lack of decision-procedure'. The idea is that consequentialists in general, and utilitarians in particular, do

not have a procedure that allows individuals to guide their actions according to the actual consequences.

Lenman considers several potential responses to these worries, but rejects them. Would you endorse any of these responses? What are these responses? Do you think they can be strengthened (do you think that Lenman has considered them in their strongest form)? One suggested response is to distinguish between subjective and objective senses of rightness? What is this distinction? Do you think it is helpful and do you think it is correct? Kagan defends consequentialism by considering the fact that no moral theory assumes individuals know all of the facts ("Tu Quoque") so all theories which introduce consequences face the same problem? Do you agree?

- Recommended – Joanna Burch-Brown – “Clues for Consequentialists” (2014)
  - In this recommended reading, Joanna Burch-Brown offers clues for consequentialists: in particular, if we assume the consequences lie in a traditional symmetrical bell distribution, then the positive and the negative potential effects will balance out over the long run. One of the strengths of this paper is the introduction of an important but surprising result from another field and applying it to a fundamental moral question.

What is the Central Limit Theorem? Why does Burch-Brown think that it helps solve Lenman's objections? Under which conditions does CLT hold? Do we have reason to think it holds in some, many, or all situations of moral interest?

- Assignment: Lenman
- SMBC – The Plan: <https://www.smbc-comics.com/index.php?id=1192>

### Objection 2: Can Consequentialists Respect Persons or Integrity? (23.06.2016 / Nico Schmidt)

- Bernard Williams – “Against Utilitarianism”
- Assignment – Bernard Williams
- OUP – Williams Quiz:  
<https://global.oup.com/us/companion.websites/9780199797264/student/part2/quizzes/williams/>
- SMBC – Explaining Morality: <https://www.smbc-comics.com/index.php?id=2992>

### Objection 3: Does Consequentialism Give Us Absurd Intertemporal Results? (30.06.2016 / Annelise Tate)

- Williams – “The Possibility of an Ongoing Moral Catastrophe” (2015)
  - In this very recent essay, Evan Williams gives two arguments that we are living through a moral catastrophe. The first points to past injustices that were perpetrated without people realizing their importance. The second points to the very many ways that people could get things morally wrong and the improbability that we happen to get things right.

When considering past injustices, how do you think people felt while they were living through them? Does this seem similar to current potential injustices? Does Williams' disjunctive worry describe ethical outcomes in a way that you agree with? Why or why not? Do Williams' solutions about learning or hedging against potential moral wrongs seem right to you? What do you think are the likelihoods that we could discover that we

are in a moral catastrophe through the methods he suggests? Can you think of better ones?

- Recommended – Nick Bostrom – “Infinite Ethics” (2011)
  - One fascinating question in modern consequentialism is whether consequentialism gets paralyzed by various forms of infinity. One of these forms is the potential (or actual?) unendingness of the universe. In this essay, Bostrom discusses some of the problems this would have for consequentialism.

Should the size of the universe affect ethics? Why would it affect consequentialist ethics differently from non-consequentialist ethics? When you consider a potentially infinite universe, does that make your ethical choices matter less, more, or the same? Why or why not? Are any of the potential solutions to this problem that Bostrom considers convincing to you? Do you have other potential solutions?

- Assignment: Williams
- SMBC – Work Experience: <https://www.smbc-comics.com/index.php?id=3553>
- SMBC – Utilitarian Computer: <https://www.smbc-comics.com/?id=2569>
- On Writing a Philosophy Paper:  
<http://people.hss.caltech.edu/~franz/Knowledge%20and%20Reality/How%20to%20Write%20a%20Philosophy%20Paper.htm>
  - You should be able to find helpful information to help you approach and structure your final paper in the links here.