

Utilitarianism and Practical Ethics

Are we to extend our concern to all the beings capable of pleasure and pain whose feelings are affected by our conduct? or are we to confine our view to human happiness? The former view is the one adopted by... the Utilitarian school... it seems arbitrary and unreasonable to exclude from the end, as so conceived, any pleasure of any sentient being.

– [Henry Sidgwick](#)¹

Table of Contents

1. Introduction
2. Is There a Difference Between Doing and Allowing Harm?
3. Demandingness
4. Cause Impartiality
5. The Expanding Moral Circle
 - Cosmopolitanism: Expanding the Moral Circle Across Geography
 - Anti-Speciesism: Expanding the Moral Circle Across Species
 - Longtermism: Expanding the Moral Circle Across Time
6. Respecting Commonsense Moral Norms
7. Conclusion
8. Resources and Further Reading

Introduction

Utilitarianism has important implications for how we should think about leading an ethical life. In this chapter, we focus on five of its theoretical implications. First, unlike many other ethical theories, utilitarianism does not regard actions and omissions as morally different. Second, it is unusually demanding: it asks us to sacrifice more than many other ethical theories do. Third, it implies that we should be *cause-impartial*: that we should apply our altruistic efforts wherever we can have the most positive impact. Fourth, it urges us to consider the well-being of individuals regardless of what country they live in, what species they belong to, and at what point in time they exist. Fifth, despite differing radically from commonsense morality as an approach to ethics, utilitarianism generally does endorse commonsense moral prohibitions.

We discuss how utilitarians should conduct their lives in the article [Acting on Utilitarianism](#). In brief, most utilitarians should [donate a significant portion of their income](#) to address the world's most pressing problems, [devote their careers to doing good](#), and aspire to high degrees of cooperativeness, personal integrity, and honesty.

Is There a Difference Between Doing and Allowing Harm?

“A person may cause evil to others not only by his actions but by his inaction, and in either case he is justly accountable to them for the injury.”

– [John Stuart Mill](#), [On Liberty](#)

Many non-consequentialists believe there is a morally relevant difference between [doing harm and allowing harm](#), even if the consequences of an action or inaction are the same. This position is known as the Doctrine of Doing and Allowing, according to which harms caused by actions—by things we actively do—are worse than harms of omission. Those who subscribe to this doctrine may, for instance, claim that it's worse to harm a child, all else being equal, than it is to fail to prevent the same child from being harmed in an accident.

Of course, all else is typically not equal. From the perspective of consequentialists and non-consequentialists alike, a societal norm allowing people to harm children would be worse than one allowing people to neglect preventing children from being harmed accidentally. This is because actively harming children could set a precedent encouraging others to harm children more, which would have worse overall consequences. Doing harm may well be *instrumentally* worse than allowing harm even if the *Doctrine of Doing and Allowing* is false.

However, while consequentialists—including utilitarians—accept that doing harm is typically instrumentally worse than allowing harm, they deny that doing harm is intrinsically worse than allowing harm. For utilitarians, only the outcomes matter. To the harmed child, the result is the same—whether you do the harming, or someone else does and you fail to prevent it. When considered from the child's perspective, the action-omission distinction is irrelevant: whether their suffering results from your deliberate action or your neglect, they suffer the same either way.

It matters a great deal whether or not there is an intrinsic moral difference between doing and allowing harm. As pointed out above, the *Doctrine of Doing and Allowing* is at the heart of the disagreement between many consequentialists and non-consequentialists. Furthermore, it matters for real-world decision-making. For instance, the ethics guidelines of many medical associations never allow doctors to actively and intentionally cause the death of a patient. However, it's acceptable for doctors to intentionally let a patient die under certain circumstances, such as if the patient is in great pain and wants to end their life. This distinction—between a doctor letting a terminally ill patient die and a doctor actively assisting a patient who wants to end their life—is regarded as less relevant from a utilitarian perspective. If we allow doctors to let a terminally ill

patient (who wishes to end their suffering) die, a utilitarian would argue that doctors should also be permitted to actively assist a patient to bring about their death with their consent.

Demandingness

Utilitarianism is typically understood to be a very [demanding](#) ethical theory: it maintains that any time you can do more to help other people than you can to help yourself, you should do so.² For example, if you could sacrifice your life to save the lives of several other people then, other things being equal, according to utilitarianism, you ought to do so.

Though occasions where sacrificing your own life is the best thing to do are rare, utilitarianism may realistically recommend many other significant sacrifices. For example, by [donating to a highly effective global health charity](#), you can save a child's life for just a few thousand dollars.³ As long as the benefits others receive from such donations exceeds what you would gain from keeping the money for yourself—as they almost certainly do, if you are a typical citizen of an affluent country—you should give it away. Indeed, you should probably donate the majority of your lifetime income. According to utilitarianism, it's only truly justifiable to spend money on yourself—such as by going out to the movies or buying nice clothes—if you think that this expenditure would do more good than any possible donation (for example by helping you work harder so that you will subsequently give away even more). This is a very high bar.

As well as recommending very significant donations, utilitarianism claims that you ought to [choose whatever career will most benefit others](#), too. This might involve non-profit work, conducting important research, or going into politics or advocacy.

Cause Impartiality

The prevailing view on helping others is that whom we should help and how we should help them is a matter of personal preference. On this common view, one may choose whether to focus on education, the arts, endangered species, or some other cause on the basis of one's personal passions.⁴

However, utilitarians argue that we should not choose our focus based primarily on personal attachment to a social cause; instead, we should apply our focus wherever we can have the most positive impact on others. Utilitarianism entails what we may call *cause impartiality*.⁵

Cause impartiality is the view that one's choice of social cause to focus on should just depend on the expected amount of good that one would do by focusing on that cause.

To illustrate this idea, suppose that you could donate to one of two different charities. One provides bednets to protect children from malaria; the other treats cancer patients. Suppose that you will save twice as many lives by donating to the bednet charity than by donating to the cancer charity;

however, a family member died of cancer, so you have a personal attachment to that cause. Should you, therefore, give to the cancer charity?

The utilitarian would argue not. On the utilitarian view, the importance of saving twice as many lives outweighs the personal attachment that might bring the donor to prioritize the cancer sufferers. From the viewpoint of utilitarianism, we should be completely impartial in deciding which social cause to support, and instead let this decision be driven only by the question of where we can do the most good.

Importantly, we know that some ways of benefiting individuals do much more good than others. For example, within the cause of global health and development, some interventions are over 100 times as effective as others.⁶ Furthermore, many researchers believe that the difference in expected impact among *causes* is as great as the differences among *interventions within a particular cause*. If so, focusing on the very best causes is vastly more impactful than focusing on average ones.

The Expanding Moral Circle

A crucial question in deciding which cause to dedicate our efforts to regards which individuals we should include in our moral deliberations.

We now recognize that characteristics like race, gender, and sexual orientation do not justify discriminating against individuals or disregarding their suffering. Over time, our society has gradually expanded our moral concern to ever more groups, a trend of moral progress often called the *expanding moral circle*.⁷ But what are the limits of this trend? Should the moral circle include all humans but stop there? Should it be expanded to include non-human animals as well? Or should it ultimately extend even to plants and the natural environment?

Utilitarianism provides a clear response to this question: We should extend our moral concern to all *sentient beings*, meaning every individual capable of experiencing positive or negative conscious states. This includes humans and many non-human animals, but presumably not plants or other non-sentient entities. This view is sometimes called *sentiocentrism* as it regards sentience as the characteristic that entitles individuals to moral concern. Justifying this perspective, Peter Singer writes:

The capacity for suffering and enjoying things is a prerequisite for having interests at all, a condition that must be satisfied before we can speak of interests in any meaningful way. It would be nonsense to say that it was not in the interests of a stone to be kicked along the road by a child. A stone does not have interests because it cannot suffer. Nothing that we can do to it could possibly make any difference to its welfare. A mouse, on the other hand, does have an interest in not being tormented, because mice will suffer if they are treated in this way...

If a being suffers, there can be no moral justification for refusing to take that suffering into consideration. No matter what the nature of the being, the principle of equality requires that the suffering be counted equally with the like suffering—in so far as rough comparisons can be made—of any other being. If a being is not capable of suffering, or of experiencing enjoyment or happiness, there is nothing to be taken into account. This is why the limit of sentience... is the only defensible boundary of concern for the interests of others.⁸

On this basis, a priority for utilitarians may be to help society to continue to widen its moral circle of concern. For instance, we may want to persuade people that they should help not just those in their own country, but also those on the other side of the world; not just those of their own species but all sentient creatures; and not just people currently alive but any people whose lives they can affect, including those in generations to come.

Cosmopolitanism: Expanding the Moral Circle Across Geography

According to utilitarianism, geographical distance and national membership are not intrinsically morally relevant. This means that, by the lights of utilitarianism, we have no grounds for discriminating against someone because of where they live, where they come from, or their nationality. This makes utilitarianism an example of [*moral cosmopolitanism*](#). Proponents of moral cosmopolitanism believe that if you have the means to save a life in a faraway country, doing so matters just as much as saving a life close by in your own country; all lives deserve equal moral consideration, wherever they are.

Of course, the geographical distance between oneself and one's beneficiary may matter instrumentally—it's often easier to help people close by than people far away. However, in an increasingly globalized world it has become much easier to benefit even those who live on the other side of the world. And because of extreme global economic inequalities, an additional unit of resources benefits people in the least-developed countries much more than people in affluent countries like the United States or the United Kingdom—potentially 100 to 1,000 times more.⁹

We discuss the implications of cosmopolitanism for ethical action in the article [Acting on Utilitarianism](#).

Anti-Speciesism: Expanding the Moral Circle Across Species

Utilitarianism cares not only about the well-being of humans, but also about the well-being of non-human animals. Consequently, utilitarianism rejects [*speciesism*](#): the practice of giving some individuals less moral consideration than others or treating them worse based on their species membership. To give individuals moral consideration is simply to consider how one's behavior will affect them, whether by action or omission. As [Peter Singer](#) describes it:

Racists violate the principle of equality by giving greater weight to the interests of members of their own race when there is a clash between their interests and the interests of those of

another race. Sexists violate the principle of equality by favoring the interests of their own sex. Similarly, speciesists allow the interests of their own species to override the greater interests of members of other species. The pattern is identical in each case.¹⁰

There is a growing scientific consensus that many non-human animals are sentient,¹¹ though not necessarily to the same degree. This includes most vertebrates, such as mammals, birds and fish, and potentially some invertebrates, such as octopodes or even insects. These animals can feel pleasure and pain, and these experiences are morally relevant from a utilitarian perspective.

Rejecting speciesism entails giving *equal moral consideration* to the well-being of all individuals, but does not entail treating all species equally. Species membership is not morally relevant *in itself*, but individuals belonging to different species may differ in other ways that do matter morally. In particular, it's likely that individuals from different species do not have the same capacity for conscious experience—for instance, because of the differing numbers of neurons in their brains. Since utilitarians believe that only sentience matters morally in itself, the utilitarian concern for individuals is proportional to their capacity for conscious experience. It's perfectly consistent with a rejection of speciesism to say we should equally consider the well-being of a fish and a chimpanzee, without implying that they have the capacity to suffer to the same degree.

From the utilitarian perspective, what matters *intrinsically* is the well-being of individual sentient beings, not the survival of the species, or the integrity of the ecosystem, or of nature. Individuals can suffer, while a “species”, an “ecosystem”, or “nature” cannot. Of course, the survival of species and the integrity of ecosystems and of nature may well be important instrumentally, to the extent that they contribute to the well-being of individuals.

Speciesism underlies the current exploitation of billions of non-human animals by humans. Animals are widely seen as resources: raised and slaughtered for food, used for clothing, or exploited for their work. These practices often result in the animals experiencing extreme suffering.

However, not all animal suffering is caused by humans. There are many more wild animals living in nature than there are domesticated animals.¹² In contrast to the widespread romanticized view of nature, wild animals generally live short lives in harsh environments. They experience suffering from many sources including predation, disease, parasites, exposure to extreme heat or cold, hunger, thirst, and malnutrition. Against this background, it would be wrong to consider only the well-being of domesticated animals which humans actively harm, while ignoring the well-being of wild animals which humans merely allow to be harmed.¹³ As noted earlier, for the utilitarian, the distinction between doing and allowing harm is irrelevant. Therefore, from the utilitarian viewpoint, we should care equally about the welfare of domestic and wild animals. That said, we currently know little about how to systematically improve the lives of wild animals. By contrast, reducing society's consumption of factory-farmed meat, or improving conditions on factory farms, would yield clear and enormous benefits for animals.¹⁴

We discuss the implications of rejecting speciesism for ethical action in the article [Acting on Utilitarianism](#).

Longtermism: Expanding the Moral Circle Across Time

From the utilitarian perspective, people on the other side of the planet matter no less than people closer to us geographically. In the same way, utilitarianism regards the well-being of future generations as no less important simply because they are far away in time than the well-being of those alive today.

A striking fact about the history of civilization is just how early in it we appear to be. There are 5,000 years of recorded history behind us, but how many years are potentially still to come? If we merely last as long as the typical mammalian species, we still have 200,000 years to go. But humans are not typical mammals, and if we can preserve our species, there could be a further one billion years until the Earth is no longer habitable,¹⁵ and other planets and solar systems will be around for trillions of years. Even on the most conservative of these timelines, we've progressed through a tiny fraction of recorded history. If humanity's saga were a novel, we would still be on the first page.

There could be astronomically more people in the future than in the present generation. This strongly suggests that, to help people in general, your key concern should not be to merely help the present generation, but to ensure that the long-term future goes as well as possible.¹⁶ This idea is known as *strong longtermism*:

Strong longtermism is the view that the most important determinant of the value of our actions today is how those actions affect the very long-run future.

Strong longtermism is implied by most plausible forms of utilitarianism¹⁷ if we assume that some of our actions can meaningfully affect the long-term future and that we can estimate which effects are positive and which negative. For example, there are risks to the continued survival of the human race, including from nuclear war, extreme climate change, man-made pathogens, and artificial general intelligence.¹⁸ If we believe that the continued survival of the human race is positive in value, then reducing the risk of human extinction is a way of positively influencing the very long-run future.¹⁹ A discussion of longtermism would go beyond the scope of this chapter, but to learn more, we recommend reading [this academic paper](#).²⁰

We discuss the implications of longtermism for ethical action in the article [Acting on Utilitarianism](#).

Respecting Commonsense Moral Norms

Sometimes people mistake utilitarianism as claiming that *one ought always to explicitly calculate the expected value of each possible action, and do whatever act scores highest*. Utilitarianism does *not* in fact recommend adopting this “naïve utilitarian” decision procedure.²¹

Instead, as a [multi-level theory](#), utilitarianism specifies moral goals—criteria for objectively judging the moral merits of an action, given all the relevant factual details²²—but leaves open the question of what kind of *decision procedure* we should try to follow in practice. After all, it’s an open empirical question how best to actually achieve the specified moral goals.

Utilitarianism implies that we should use whatever decision procedure would best help us to promote overall well-being (in expectation). While we cannot be certain what decision procedure satisfies this criterion, we can offer some educated guesses. As psychologists Stefan Schubert and Lucius Caviola argue in [Virtues for Real-World Utilitarians](#), we may best promote overall well-being by ambitiously pursuing robustly good [actions that effectively help others](#), while minimizing downside risks by means of commonsense virtues and constraints.

It’s widely recognized that humans are unreliable at calculating utilities,²³ especially when they conflict with generally-reliable rules and heuristics (such as those prohibiting harm to others). As a result, we cannot take rule-violating expected value calculations at face value. Even if you’ve calculated that it would somehow serve the greater good to murder your rival, you should be very skeptical that this is true. After all, if you don’t really believe that it would overall be best for others (similarly situated) to do likewise, then you must believe that most calculations favoring murder have actually gone awry. So, if you’ve no special (symmetry-breaking) evidence establishing that you, in particular, are the rare exception to this rule, then you must conclude that your own murderous calculations have most likely gone awry.²⁴ Thus, absent special evidence, you should conclude that your rule-breaking actually has lower expected value, despite your initial calculation to the contrary.

We can be most confident that our actions have positive expected value when we instead seek to help others in ways that are supported by good evidence and minimize downside risk.²⁵ Over the long run, we should expect honest, cooperative altruism to do more good than ruthless scheming for the “greater good”, because (i) historically, ruthless schemers often do more harm than good, (ii) people rightly don’t trust ruthless schemers, and (iii) in a complex world, it’s difficult to get much done without others’ trust and cooperation. If that’s right, honest cooperative altruism systematically has *higher expected value* than ruthless scheming, and should be preferred by utilitarians.

In summary, utilitarianism does *not* tell us to constantly calculate utilities and blindly follow whatever our calculations recommend. That would be predictably counter-productive, contrary to the pragmatic spirit of the theory. Instead, utilitarianism recommends decision-procedures based on their expected value. When we are uncertain, we should be guided by whatever decision

procedure can most reasonably (in light of everything we know about human biases and cognitive limitations) be expected to yield better outcomes. This means following *heuristics* or generally-reliable rules of thumb.

As a very rough first pass, a plausible utilitarian decision-procedure might direct us to:²⁶

1. Pursue any “low-hanging fruit” for [effectively helping others](#) while avoiding harm,
2. Inculcate [virtues for real-world utilitarians](#) (including respect for commonsense moral norms), and
3. In a calm moment, reflect on how we could better prioritize and allocate our moral efforts, including by seeking out expert cost-benefit analyses and other evidence to better inform our overall judgments of expected value.²⁷

Conclusion

Utilitarianism has important implications for how we should think about leading an ethical life.

The theory rejects an intrinsic moral difference between doing and allowing harm. This position contributes to the demandingness of utilitarianism, since it implies that whenever we decide not to help another person, we are complicit in their misery.

By the lights of utilitarianism, we should choose carefully which moral problems to work on and by what means, based on where we can do the most good. We should extend our moral concern to all sentient beings, meaning every individual capable of experiencing happiness or suffering.

Utilitarianism urges us to consider the well-being of all individuals regardless of what species they belong to, what country they live in, and at what point in time they exist.

Though utilitarians should try to use their lives to do the most good they can, in practice, they should do so while respecting commonsense moral virtues like honesty, integrity, fairness, and law-abidingness. There are reasons we do not see utilitarians robbing banks to donate the proceeds: these commonsense moral prohibitions help society to function smoothly, and any naive calculation that violating such a prohibition would promote the greater good is almost always mistaken.

The next chapter discusses important rival theories that may overlap significantly with utilitarianism in practice.

Next Chapter: Near-Utilitarian Alternatives

Como Citar esta Página

MacAskill, W., Meissner, D., and Chappell, R.Y. (2023). Utilitarianism and Practical Ethics. In R.Y. Chappell, D. Meissner e W. MacAskill (eds.), *An Introduction to Utilitarianism*, <<https://www.utilitarismo.net/utilitarianism-and-practical-ethics>>, acessado em 23/05/2025.

Resources and Further Reading

Is There a Difference Between Doing and Allowing Harm?

- Fiona Woollard & Frances Howard-Snyder (2016). [Doing vs. Allowing Harm](#). *The Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.).
- Jonathan Bennett (1995). *The Act Itself*. Oxford University Press.

The Expanding Moral Circle

- Peter Singer (1997). [The Drowning Child and the Expanding Circle](#). *New Internationalist*.
- Peter Singer (2011). [The Expanding Circle: Ethics, Evolution, and Moral Progress](#). Princeton: Princeton University Press.

Cosmopolitanism: Expanding the Moral Circle Across Geography

- [Poverty & Our Response to It: Crash Course Philosophy #44](#)
- Peter Singer (2019). [The Life You Can Save: Acting Now to End World Poverty](#), 2nd ed. The Life You Can Save, Bainbridge Island, WA and Sydney, available free at <www.thelifeyoucansave.org>.
- Peter Singer (1972). [Famine, Affluence, and Morality](#). *Philosophy & Public Affairs*. 1(2): 229–243.
- Samuel Scheffler (1999). [Conceptions of Cosmopolitanism](#). *Utilitas*. 11(3): 255–276.

Anti-Speciesism: Expanding the Moral Circle Across Species

- [Non-Human Animals: Crash Course Philosophy #42](#)
- Peter Singer (2023) *Animal Liberation Now: The Definitive Classic Renewed*, New York: HarperCollins.
- Jeff McMahan (2002). Animals. In R. G. Frey and Christopher Wellman (eds.), *The Blackwell Companion to Applied Ethics*. Oxford: Blackwell, pp. 525–536.
- Jeff Sebo (2019). [A Utilitarian Case for Animal Rights](#). *Effective Altruism Global*.

Longtermism: Expanding the Moral Circle Across Time

- Toby Ord (2020). [*The Precipice: Existential Risk and the Future of Humanity*](#). London: Bloomsbury Publishing.
- Hilary Greaves & William MacAskill (2019). [*The Case for Strong Longtermism*](#). *Global Priorities Institute*.
- Nick Beckstead (2013). [*On the Overwhelming Importance of Shaping the Far-Future*](#). PhD Dissertation, Rutgers University.
- Nick Bostrom (2003). [*Astronomical Waste: The Opportunity Cost of Delayed Technological Development*](#). *Utilitas*. 15(3): 308–314.
- William MacAskill (2022). [*What We Owe the Future*](#). Basic Books.

Respecting Commonsense Moral Norms

- Allan Gibbard (1984). [*Utilitarianism and Human Rights*](#). *Social Philosophy and Policy*, 1(2): 92–102.
- R.M. Hare (1981). [*Moral Thinking*](#). Oxford University Press.
- J.L. Mackie (1985). Rights, Utility, and Universalization. In R.G. Frey (ed.) *Utility and Rights*. Basil Blackwell.
- Philip Pettit & Geoffrey Brennan (1986). [*Restrictive Consequentialism*](#). *Australasian Journal of Philosophy*, 64(4): 438–455.

-
1. Sidgwick, H. (1981). [*The Methods of Ethics*](#). Hackett Publishing. Book IV, p. 414 ↻
 2. Though as flagged in our chapter on [the demandingness objection](#), this may not actually be so demanding if the “should” claim merely indicates what is morally *ideal*, rather than what is *required to avoid deserving moral blame or criticism*.

For a discussion of demandingness in the context of global poverty alleviation, see Singer, P. (2019). [*The Life You Can Save: Acting Now to End World Poverty*](#), 2nd ed. The Life You Can Save, Bainbridge Island, WA and Sydney, available free at <www.thelifeyoucansave.org>. ↻
 3. GiveWell (2019). [*Your dollar goes further overseas*](#). ↻
 4. Berman, J. Z., Barasch, A., Levine, E. E., & Small, D. A. (2018). [*Impediments to Effective Altruism: The Role of Subjective Preferences in Charitable Giving*](#). *Psychological Science*. 29(5): 834–844. ↻
 5. The case for cause neutrality is made in Effective Altruism Foundation (2017). [*The Benefits of Cause-Neutrality*](#). ↻

6. Ord, T. (2013). [The Moral Imperative Towards Cost-Effectiveness in Global Health](#). *Center for Global Development*. ↩
7. Cf. Singer, P. (2011). [The Expanding Circle: Ethics, Evolution, and Moral Progress](#). Princeton: Princeton University Press. ↩
8. Singer, P. (2011). [Practical Ethics](#). Cambridge: Cambridge University Press, p. 50 ↩
9. Cf. MacAskill, W. (2015). [Doing Good Better: Effective Altruism and How You Can Make a Difference](#). New York: Penguin Random House. Chapter 1. ↩
10. Singer, P. (2023) [Animal Liberation Now: The Definitive Classic Renewed](#), New York: HarperCollins, p. 8. Indeed, there is psychological evidence suggesting that speciesism goes hand in hand with other discriminatory attitudes like racism, sexism, and homophobia: Cf. Caviola, L; Everett, J. A. C. & Faber, N. S. (2017). [The Moral Standing of Animals: Towards a Psychology of Speciesism](#). *Journal of Personality and Social Psychology*. 116(6): 1011–1029. ↩
11. For instance, see the [Cambridge Declaration on Consciousness](#) from Low, P. et al. (2012) ↩
12. Tomasik, B. (2019). [How Many Wild Animals Are There?](#). ↩
13. There is an ongoing academic debate about the moral importance of wild animal welfare. For example, see the following: Ng, Y. (1995). [Towards Welfare Biology: Evolutionary Economics of Animal Consciousness and Suffering](#). *Biology and Philosophy*. 10 (3): 255–285. McMahan, J. (2013). [The Moral Problem of Predation](#). In Chignell, A; Cuneo, T. & Halteman, M. (eds.). *Philosophy Comes to Dinner: Arguments on the Ethics of Eating*. London: Routledge. Moen, O. M. (2016). [The Ethics of Wild Animal Suffering](#). *Etikk I Praksis - Nordic Journal of Applied Ethics*. 10: 1–14. ↩
14. For a further discussion of this topic we recommend this interview with researcher Persis Eskander: Eskander, P. (2019). [Animals in the wild often suffer a great deal. What, if anything, should we do about that?](#). *80,000 Hours Podcast with Rob Wiblin*. ↩
15. Adams, F. C. (2008). Long-Term Astrophysical Processes. In Bostrom, N. and Cirkovic, M. (eds.) [Global Catastrophic Risks](#). Oxford: Oxford University Press. ↩
16. For a discussion of this idea and its underlying assumptions, see Beckstead, N. (2013). [On the Overwhelming Importance of Shaping the Far-Future](#). PhD Dissertation, Rutgers University. ↩
17. Cf. Greaves, H. & MacAskill, W. (2019). [The Case for Strong Longtermism](#). *Global Priorities Institute*. Section 4.1. ↩
18. For a detailed discussion of existential risks and the moral importance of the long-run future of humanity, see Ord, T. (2020). [The Precipice: Existential Risk and the Future of Humanity](#).

London: Bloomsbury Publishing. ↩

19. For a classic paper on the importance of reducing existential risk, see Bostrom, N. (2013). [Existential Risk Prevention as Global Priority](#). *Global Policy*. 4 (1): 15–31. ↩
20. Note that William MacAskill, coauthor of this website, is also a coauthor of this paper. ↩
21. As [Chapter 2 explains](#): “[T]o our knowledge no one has ever defended single-level utilitarianism [i.e., using the utilitarian criterion as a universal decision procedure], including the classical utilitarians. Deliberately calculating the expected consequences of our actions is error-prone and risks falling into decision paralysis.” ↩
22. Such details might simply be stipulated in a hypothetical example. In real life cases, our uncertainty about relevant factual details should generally carry over to make us similarly uncertain about our moral verdicts and evaluations. ↩
23. See, e.g., J.L. Mackie (1985). Rights, Utility, and Universalization. In R.G. Frey (ed.) *Utility and Rights*. Basil Blackwell. ↩
24. In particular, you can’t take at face value your inclination to think that there are special reasons in your case, if you believe that most people in subjectively similar situations are mistaken in taking themselves to have such special reasons. Symmetry-breaking evidence is evidence that *distinctively* establishes your reliability in comparison to others with similar (but, in their case, misguided) beliefs. Note that such symmetry-breaking evidence is very hard to come by! ↩
25. That’s not to say that we should strictly optimize for *confidence* in positive expected value (EV): Something that’s *certainly barely-good* in expectation may be less worth pursuing than something that is *almost* certainly high EV even if there’s a slight risk that you’ve overlooked something that would mean the action is actually mildly negative in expectation. Such uncertainty could still result in higher “all things considered” expected value, in principle. The point is just that grounds for doubting a positive verdict from our initial EV calculations should rationally lead us to assign lower (and in some cases, even *negative*) expected value to that option, all things considered. And in practice, it seems that we should often have strong priors that rights-violating actions are net-negative, which a rough and broadly unreliable calculation should not suffice to overturn. ↩
26. These claims are not, strictly speaking, built into utilitarianism as a fundamental moral theory. Rather, we are speculating about the further question of *what decision-procedure has the highest expected value in typical circumstances*. (Note that, in principle, the answer may differ for individuals in different circumstances. Nothing in utilitarianism requires uniformity, if that would not be for the best.) ↩

27. This might (but need not) include performing some “back of the envelope” calculations of expected value. Even then, to truly maximize expected value, these naive calculations must be tempered by constraints against ruthless scheming, given our prior judgment that the latter is most likely counterproductive. That is, if we calculate a slightly higher explicit “expected value” for one act than another, but the former involves egregious norm-breaking, we should probably conclude that the latter (safer) option is *actually* better in expectation. ↩