

Glossary

This page provides an overview with brief descriptions of key utilitarian terms and relevant links.

[Act utilitarianism](#) ^

Act utilitarianism is the view that one morally ought to promote just the sum total of well-being.

¹ Act utilitarianism is the best known version of [direct consequentialism](#) and is often contrasted with *rule utilitarianism*, an indirect consequentialist view. Contemporary utilitarian philosophers often endorse [global utilitarianism](#), which emphasizes that utilitarian standards of moral evaluation apply to anything of interest (not just acts). ²

[Aggregationism](#) ^

→ *Main article: [Aggregationism](#)*

Aggregationism holds that the value of the world is the sum of the values of its parts, where these parts are local phenomena such as experiences, lives, or societies. ³ When combined with welfarism and the equal consideration of interests, this view implies that we can meaningfully add up the well-being of different individuals, and use this total to determine which trade-offs are worth making. Aggregationism is one of the [four elements of utilitarian ethical theories](#).

[Arguments in favor of utilitarianism: theoretical virtues](#) ^

→ *Main article: [Arguments for utilitarianism](#)*

Utilitarianism has strong theoretical virtues as an ethical theory. It is simple and clear, and it provides concrete implications for how to act in any situation.

[Arguments in favor of utilitarianism: track record](#) ^

→ *Main article: [Arguments in favor of utilitarianism: Track record](#)*

Utilitarian moral reasoning has a strong track record of contributing to humanity's collective moral progress. The classical utilitarians of the 18th and 19th centuries—[Jeremy Bentham](#), [John Stuart Mill](#), and [Henry Sidgwick](#)—had social and political attitudes that were far ahead of their time. While the early proponents of utilitarianism were still far from getting everything right, their utilitarian reasoning led them to escape many of their time's moral prejudices and develop more enlightened moral views. Utilitarianism enabled Bentham, Mill, and Sidgwick to make

better moral “predictions” than those who endorsed alternative moral views. That is, utilitarianism led the early utilitarians to many conclusions which struck people as counterintuitive at the time but which most of us now understand as right. This provides us with some reason to expect that when today’s “common sense” moral intuitions conflict with utilitarian conclusions, the latter are more likely to be correct. At the very least, checking our moral and political views against utilitarian principles may help us to avoid and overcome some of our own biases.

Arguments in favor of utilitarianism: the veil of ignorance

→ *Main article: [Arguments for utilitarianism: The Golden Rule, the Veil of Ignorance, and the Ideal Observer](#)*

Imagine you had to decide how to structure society from behind a [veil of ignorance](#). Behind this veil of ignorance, you know all the facts about each person’s circumstances in society—what their income is, how happy they are, how they are affected by social policies, and their preferences and likes. However, what you do not know is which of these people you are. You only know that you have an *equal chance* of being any of these people. Imagine, now, that you are trying to act in a rational and self-interested way—you are just trying to do whatever is best for yourself. How would you structure society?

Nobel Prize-winning economist John Harsanyi proved that in this situation you will structure society to promote the sum total of everyone’s well-being.⁴ In other words, if you are rational and acting in self-interest and were put behind the veil of ignorance, you would come to use some version of utilitarianism as the principle to decide about the structure and rules of society.

Astronomical waste

Oxford philosopher Nick Bostrom [writes that](#) “With very advanced technology, a very large population of people living happy lives could be sustained in the accessible region of the universe. For every year that development of such technologies and colonization of the universe is delayed, there is therefore a corresponding opportunity cost: a potential good, lives worth living, is not being realized”.⁵ He coined the term “astronomical waste” to describe this opportunity cost of delayed technological development. Bostrom argues that, despite this large opportunity cost, utilitarians should not aim to maximize the rate of technological progress “but rather that we ought to maximize its safety, i.e. the probability that colonization will eventually occur”.⁶

See also: [Existential risk reduction](#)

Average view (population ethics)

→ *Main article: [Average view \(population ethics\)](#)*

The average view of population ethics regards one outcome as better than another if and only if it contains greater average well-being. Since the average view aims only to improve the average well-being level, it disregards—in contrast to [the total view](#)—the number of individuals that exist. The average view avoids the [repugnant conclusion](#), because it states that reductions in the average well-being level can never be compensated for by adding more people to the population.

However, the average view has very little support among moral philosophers, because it leads to counterintuitive implications which are said to be at least as serious as the repugnant conclusion.⁷ For instance, drawing on the work of Derek Parfit⁸, Gustaf Arrhenius et al. (2017) writes that the average view implies the following: “[F]or a population consisting of just one person leading a life at a very negative level of well-being, e.g., a life of constant torture, there is another population which is better even though it contains millions of lives at just a slightly less negative level of well-being”.⁹

The main alternatives to the average view of population ethics are the [total view](#) and [person-affecting views](#). According to the total view, one outcome is better than another if and only if it contains a greater sum total of well-being, even if that is in virtue of simply having more people. Person-affecting views are a family of views that share the intuition that an act can only be good or bad if it is good or bad *for* someone. Standard person-affecting views stand in opposition to the total view, since they entail that there is no moral good in bringing new people into existence because nonexistence means there is no one for whom it could be good to be created.

Career choice

→ *Main article: [Career choice](#)*

Most of us will spend around 80,000 hours during our lives on our professional careers, and some careers achieve much more good than others. Your choice of career is, therefore, one of the most important moral choices of your life. By using this time to address the most pressing global problems, we can do an enormous amount of good. Yet, it is far from obvious which careers will allow you to do the most good from a utilitarian perspective.

Fortunately, there is research available to help us make more informed choices. The organization [80,000 Hours](#)¹⁰ aims to help people use their careers to solve the world’s most pressing problems. To do this, they research how individuals can maximize the social impact of their careers, create online advice, and support readers who might enter priority areas.

Cause impartiality and cause prioritization

[Cause impartiality](#) is the view that one’s choice of social cause to focus on should depend on, and only on, the expected amount of good that one can do in that cause. Which causes will allow us

to do the greatest amount of good by promoting well-being? Finding the answer to that question is called [cause prioritization](#).

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.¹¹ 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.

Charitable giving

→ *Main article:* [Charitable giving](#)

In slogan form, the utilitarian recommendation for using your money to help others is to “give more and give better”. Giving more simply means increasing the proportion of your income you give to charity. Giving better means finding and donating to the organizations that make the best use of your donation.

Citizens of affluent countries are in the richest few percent of the world’s population. By making small sacrifices, those in the affluent world have the power to dramatically improve the lives of others. Due to the extreme inequalities in wealth and income, one can do a lot more good by giving money to those most in need than by spending it on oneself.¹²

To give better, one can follow the recommendations from organizations such as [GiveWell](#), which conducts exceptionally in-depth charity evaluations. GiveWell’s best-guess estimate is that the most cost-effective charities working in global health can save a child’s life for about \$5,000.¹³

Classical utilitarianism

→ *Main article:* [Classical utilitarianism](#)

Classical utilitarianism is the view that one morally ought to promote just the sum total of happiness over suffering. Classical utilitarianism can be distinguished from the wider utilitarian family of views because it accepts [hedonism](#) as a theory of well-being and [the total view](#) of population ethics.

Consequentialism

→ *Main article:* [Consequentialism](#)

Consequentialism is the view that the moral rightness of actions (or rules, policies, etc.) depends on, and only on, the value of their consequences. Thus, to evaluate whether an action is right or wrong, we should look at all its consequences rather than any of its other features. For

instance, when breaking a promise has bad consequences—as it usually does—consequentialists consider it wrong to do so. However, breaking a promise is not considered wrong in and of itself. In exceptional cases breaking a promise would be morally permissible or even required, such as when doing so is necessary to save a life.

Consequentialism is one of the [four elements of utilitarian ethical theories](#).

External links: [Consequentialism](#), Stanford Encyclopedia of Philosophy

Cosmopolitanism

→ *Main article:* [Cosmopolitanism](#)

Moral cosmopolitanism is the view 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.

Utilitarianism accepts moral cosmopolitanism and consequently regards geographical distance and national membership as 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 what nationality they have.

An implication of accepting moral cosmopolitanism is to take [improving global health and development](#) very seriously as moral priorities.

External links: [Taxonomy of Contemporary Cosmopolitanisms](#), Stanford Encyclopedia of Philosophy

Demandingness

→ *Main article:* [Demandingness](#)

Utilitarianism is 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. 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 is still very demanding in the world today. For example, by [donating to a highly effective global health charity](#), you can save a child's life for just a few thousand dollars.¹⁴ As long as such donations benefit others more than a few thousand dollars would benefit yourself—as they almost certainly do, if you are a typical citizen of an affluent country—you ought to donate. Indeed, you likely ought to donate the majority of your lifetime income.

As well as requiring 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.

See also: [Demandingness Objection to Utilitarianism](#)

Demandingness objection to utilitarianism

→ *Main article:* [Demandingness objection to utilitarianism](#)

Many critics argue that utilitarianism is too demanding, because it requires us to always act such as to bring about the best outcome. The theory leaves no room for actions that are permissible yet do not bring about the best consequences; this is why some critics claim that utilitarianism is a morality only for saints.

Consider that the money a person spends on dining out could pay for several bednets, each protecting two children in a low-income country from malaria for about two years.¹⁵ From a utilitarian perspective, the benefit to the person from dining out is much smaller than the benefit to the children from not having malaria, so it would seem the person has acted wrongly in choosing to have a meal out. Analogous reasoning applies to how we use our time: the hours someone spends on social media should apparently be spent volunteering for a charity, or working harder at one's job to earn more money to donate.

See the article [The Demandingness Objection](#) on how proponents of utilitarianism might respond to this objection.

Deontology

According to *deontology*, morality is about following a system of duties and rules, like “Do Not Lie” or “Do Not Steal”. As Larry Alexander and Michael Moore [write](#): “In contrast to consequentialist theories, deontological theories judge the morality of choices by criteria different from the states of affairs those choices bring about. The most familiar forms of deontology, and also the forms presenting the greatest contrast to consequentialism, hold that some choices cannot be justified by their effects—that no matter how morally good their consequences, some choices are morally forbidden”.¹⁶

The main alternatives to deontology are [consequentialism](#), the view that the moral rightness of actions (or rules, policies, etc.) depends on, and only on, the value of their consequences, and [virtue ethics](#), according to which morality is fundamentally about having or developing a virtuous character.

External links: [Deontological Ethics](#), Stanford Encyclopedia of Philosophy

Desire theories of well-being

→ *Main article:* [Theories of well-being: Desire theories](#)

According to desire theories only the satisfaction of desires or preferences matters for an individual's well-being. The most well known desire theory is preference utilitarianism, the ethical theory on which you ought to promote just the sum total of preference satisfaction over dissatisfaction.

The alternatives to desire theories include [*hedonism*](#), according to which the individual's conscious experiences determines their well-being, and [*objective list theories*](#), which propose a list of items that constitute well-being, such as conscious experiences, art, knowledge, love, friendship, and more.

Direct consequentialism & direct utilitarianism

→ *Main article: [Consequentialism](#)*

According to direct consequentialism, the rightness of an action (or rule, policy, etc.) depends only on its consequences. On this view, to determine the right action in some set of feasible actions, we should directly evaluate the consequences of the actions to see which has the best consequences. The most well known direct consequentialist view is act utilitarianism, which assesses the moral rightness of actions, and only of actions, according to the sum total of well-being they produce.

The alternative to direct consequentialism is indirect consequentialism, according to which we should evaluate the moral status of an action (or rule, policy, etc.) *indirectly*, based on its relationship to something else (such as a rule), whose status is itself assessed in terms of its consequences.

Doctrine of Doing and Allowing

→ *Main article: [Doctrine of Doing and Allowing](#)*

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.

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. Thus, they reject the Doctrine of Doing and Allowing.

Effective altruism

→ *Main article: [Effective altruism](#)*

Those in the [effective altruism](#) movement try to figure out, of all the different uses of our resources, which ones will do the most good, impartially considered, and act on that basis. So defined, effective altruism is both a research project—to figure out how to do the most good—and a practical project to implement the best guesses we have about how to do the most good.¹⁷

Egalitarianism

→ *Main article: [Egalitarianism and Distributive Justice](#)*

Egalitarianism is the view that inequality is bad *in itself*, over and above any instrumental effects it may have on people's well-being.

Egalitarians thus reject [welfarism](#), the view that positive well-being is the only intrinsic good, and negative well-being is the only intrinsic bad.

External links: [Egalitarianism](#), Stanford Encyclopedia of Philosophy

Equal Consideration of Interests

→ *Main article: [Impartiality and the Equal Consideration of Interests](#)*

The *equal consideration of interests* is a distinctively utilitarian conception of impartiality, according to which equal weight must be given to the interests of all individuals. This means treating well-being as equally valuable regardless of when, where, or to whom it occurs.

Alternative views include [prioritarianism](#) (which gives extra weight to the interests of the worse off) and [partialism](#) (which abandons impartiality, allowing us to give extra weight to ourselves and the interests of our nearest and dearest).

Equality objection to utilitarianism

→ *Main article: [Equality objection to utilitarianism](#)*

Some argue that utilitarianism conflicts with the ideal of equality. Suppose, for example, that you could choose between two possible distributions of well-being, *Equality* and *Inequality*: Equality has 1,000 people at well-being level 45, while Inequality has 500 people at 80 well-being and another 500 people at 20 well-being.

By the lights of utilitarianism, only the sum total of well-being determines the goodness of an outcome: it does not matter how that well-being is distributed across people. Since the sum total of well-being is greater in Inequality (= 50) than in Equality (= 45), the unequal outcome is preferable according to utilitarianism. Some philosophers object to the utilitarian view regarding this choice, claiming that the equal distribution of well-being in Equality provides a reason to choose this outcome. On this view, total well-being is not all that matters; equality of

distribution also matters. Equality, it is claimed, is an important moral consideration that the utilitarian overlooks.

See the article [The Equality Objection](#) on how proponents of utilitarianism might respond to this objection.

Ex ante Pareto

→ *Main article:* [Ex Ante Pareto](#)

A *Pareto* improvement is better for some people, and worse for none. When the future is uncertain, we can assess an individual's *ax ante* interests by reference to their [expected well-being](#) (in contrast to their objective interests, which might only be knowable *ex post*, or after the fact). Putting these two concepts together, the *Ex Ante Pareto* principle holds that, in a choice between two prospects, one is morally preferable to another if it offers a better prospect for some individuals and a worse prospect for none.

(Interestingly, theories may combine *ex post* welfare evaluations with a broader “expectational” element. For example, [ex post prioritarianism](#) assigns extra social value to avoiding bad outcomes (rather than bad *prospects*) for the worst off individuals, but can still assess prospects by their *expected social value*.)

A powerful objection to many non-utilitarian views is that they are committed to violating this Ex Ante Pareto principle in some possible situations, such as when choosing policies from behind a Veil of Ignorance.

See: Harsanyi, J. C. (1955). [Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility](#). *The Journal of Political Economy*, pp. 309–321.

Existential risk reduction

→ *Main article:* [Existential risk reduction](#)

An existential risk is a risk that threatens the destruction of humanity's long-term potential—such as all-out nuclear war, or extreme climate change, or an engineered global pandemic.¹⁸ From a utilitarian perspective (and the perspective of many other moral views), the realization of an existential risk would be uniquely bad and much worse than non-existential catastrophes. Besides the deaths of all 7.8 billion people on this planet, an existential catastrophe would irreversibly deprive humanity of a potentially grand future and preclude trillions of lives to come. Since the stakes involved with existential risks are so large, their mitigation may, therefore, be one of the most important moral issues we face.

External links: [The Precipice: Existential Risk and the Future of Humanity](#), Toby Ord (2020)

Expanding moral circle

→ *Main article: [The expanding moral circle](#)*

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*.¹⁹ But what are the limits of this trend?

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 probably many non-human animals, but not plants or other entities that are non-sentient. This view is sometimes called *sentiocentrism* as it regards sentience as the characteristic that entitles individuals to moral concern.

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.

Expectational utilitarianism

→ *Main article: [Expectational utilitarianism](#)*

Expectational utilitarianism is the view we should promote *expected* well-being, as opposed to the well-being an action will *in fact* produce. Expectational utilitarianism states we should choose the actions with the highest expected value.²⁰ The expected value of an action is the sum of the value of each of the potential outcomes multiplied by the probability of that outcome occurring. So, for example, according to expectational utilitarianism, we should choose a 10% chance of saving 1,000 lives over a 50% chance of saving 150 lives because the former option saves an expected 100 lives (= 10% * 1,000 lives) whereas the latter option saves an expected 75 lives (= 50% * 150 lives).

The main alternative to expectational utilitarianism is *objective utilitarianism*, on which the rightness of an action depends on the well-being it will *in fact* produce.

Farm animal welfare

→ *Main article: [Farm animal welfare](#)*

Improving the welfare of farmed animals should be a high moral priority for utilitarians. The argument for this conclusion is simple: First, [animals matter morally](#); second, humans cause a

huge amount of unnecessary suffering to animals in factory farms; third, there are easy ways to reduce the number of farmed animals and the severity of their suffering.

Global health and development

→ *Main article: [Global health and development](#)*

Efforts in global health and development have a great track record of improving lives, making this cause appear especially tractable. Indeed, the best interventions in global health and development are incredibly cost-effective: [GiveWell](#), a leading organization that conducts in-depth charity evaluations, estimates that top-rated charities can prevent the death of a child from malaria for just a few thousand dollars by providing preventive drugs.²¹ On this basis, global health and development may be considered a particularly high priority cause for utilitarians.²²

Global utilitarianism

→ *Main article: [Global utilitarianism](#)*

Global utilitarianism is the view that the utilitarian standards of right and wrong can evaluate anything of interest, including actions, motives, rules, virtues, policies, social institutions, etc.

Global utilitarianism assesses the moral nature of, for example, a particular character trait, such as kindness or loyalty, based on the consequences that trait has for the well-being of others—just as act utilitarianism evaluates the rightness of actions. Global utilitarianism’s broad focus may help it to explain certain supposedly “non-consequentialist” intuitions.²³ For instance, it captures the understanding that morality is not just about choosing the right acts but is also about following certain rules and developing a virtuous character.

Happiness and suffering

→ *Main article: [Theories of well-being: hedonism](#)*

Philosophers commonly use *happiness* and *suffering* as shorthand for the terms *positive conscious experience* and *negative conscious experience*, respectively. According to ethical hedonists, happiness is the only thing good in and of itself and suffering is the only thing bad in and of itself. The hedonistic conception of happiness is broad: It covers not only paradigmatic instances of sensual pleasure—such as the experiences of eating delicious food or having sex—but also other positively valenced experiences, such as the experiences of solving a problem, reading a novel, or helping a friend.

Harriet Taylor Mill

→ *Main article: [Harriet Taylor Mill](#)*

Harriet Taylor Mill (1807 – 1858) was a British philosopher and women’s rights advocate. A close friend and later wife of John Stuart Mill, she had a profound impact on his thinking and worked in close collaboration with him. Despite her many contributions in books and magazines, most of her writing was only published under her own name after her death.

Hedonic calculus

[Jeremy Bentham](#) proposed the hedonic calculus, or felicific calculus, as a method to determine the goodness and badness of an action’s consequences.²⁴ Bentham suggested that in assessing these consequences, one should take into account their *intensity*, *duration*, *certainty*, *propinquity*, *fecundity* (the chance that a pleasure is followed by other ones, a pain by further pains), *purity* (the chance that pleasure is followed by pains and vice versa), and *extent* (the number of persons affected). Applying the hedonic calculus to similarly assess all the alternative actions, would show which one has the best overall consequences, and should therefore be chosen.

However, Bentham was realistic about the limitations of this method, writing that “it is not to be expected that this process [of calculating expected consequences] should be strictly pursued previously to every moral judgment”.²⁵

Hedonism (theories of well-being)

→ *Main article: [Theories of well-being: hedonism](#)*

Hedonism is the view that well-being consists in, and only in, the balance of positive over negative conscious experiences. For hedonism the only things good in and of themselves are the experiences of positive conscious states, such as enjoyment and pleasure; and the only things bad in and of themselves are the experiences of negative conscious states, such as misery and pain.

The hedonistic conception of happiness is broad: It covers not only paradigmatic instances of sensual pleasure—such as the experiences of eating delicious food or having sex—but also other positively valenced experiences, such as the experiences of solving a problem, reading a novel, or helping a friend. Hedonists claim that all these experiences are *intrinsically* valuable, which means they are valuable in and of themselves. Other goods, such as wealth, health, justice, fairness and equality are also valued by hedonists, but they are valued *instrumentally*. This means they are valued to the extent that they affect the conscious experience of individuals, rather than being valued in and of themselves.

The two main alternatives to hedonism are *desire theories*, according to which only the satisfaction of desires or preferences matters for an individual’s well-being, and *objective list*

theories, which propose a list of items that constitute well-being. This list can include conscious experiences or satisfied preferences, but it rarely stops there; ethicists commonly argue that the objective list includes art, knowledge, love, friendship, and more.

Henry Sidgwick

→ *Main article: [Henry Sidgwick](#)*

Henry Sidgwick (1838 – 1900) was a British philosopher and economist. One of the classical utilitarians, he wrote one of the most important statements of utilitarianism in his [The Methods of Ethics](#), which was said to be “the best book ever written on ethics”.²⁶

Hybrid utilitarianism

→ *Main article: [Global vs Hybrid Utilitarianism](#)*

Hybrid utilitarianism is the view that, while one morally ought to promote just overall well-being, the moral quality of an aim or intention can depend on factors other than whether it promotes overall well-being. In particular, hybrid utilitarians may understand virtue and praise-worthiness as concerning whether the target individual *intends* good results, in contrast to global utilitarian evaluation of whether the target’s intentions *produce* good results. When the two come into conflict, we should prefer to achieve good results than to merely intend them —so in this sense the hybrid utilitarian agrees with much that the global utilitarian wants to say. Hybridists just hold that there is more to say in addition.

Impartiality

→ *Main article: [Impartiality](#)*

Impartiality is the view that the identity of individuals is irrelevant to the value of an outcome. Utilitarians accept a conception of impartiality that further entails the *equal consideration of interests*: that is, the claim that equal weight must be given to the interests of all individuals. This means treating well-being as equally valuable regardless of when, where, or to whom it occurs. As a consequence, utilitarianism values the well-being of all individuals equally, regardless of their nationality, gender, [where](#) or [when they live](#), or even [their species](#).

Impartiality is one of the [four elements of utilitarian ethical theories](#).

Indirect consequentialism & indirect utilitarianism

→ *Main article: [Consequentialism](#)*

According to indirect consequentialism we should evaluate the moral status of an action *indirectly*, based on its relationship to something else (such as a rule), whose status is itself

assessed in terms of its consequences. The most well known indirect consequentialist view is rule utilitarianism, which holds that what makes an action right is that it conforms to the set of rules that would have the best utilitarian consequences if they were generally accepted or followed.

The main alternative to indirect consequentialism is direct consequentialism, according to which the rightness of an action (or rule, policy, etc.) depends only on its consequences.

Infinite ethics

In a [2011 paper](#), Nick Bostrom suggests that infinities in ethics may present a problem for aggregative consequentialist theories, including utilitarianism. Bostrom describes this problem as follows: “Modern cosmology teaches that the world might well contain an infinite number of happy and sad people and other candidate value-bearing locations. Aggregative ethics implies that such a world contains an infinite amount of positive value and an infinite amount of negative value. You can affect only a finite amount of good or bad. In standard cardinal arithmetic, an infinite quantity is unchanged by the addition or subtraction of any finite quantity. So it appears you cannot change the value of the world”.²⁷

Jeremy Bentham

→ *Main article:* [Jeremy Bentham](#)

Jeremy Bentham (1748 – 1832) was a British philosopher and social reformer, who is widely regarded as the founder of classical utilitarianism. His most influential work is [An Introduction to the Principles of Morals and Legislation](#) (1789).

John Stuart Mill

→ *Main article:* [John Stuart Mill](#)

John Stuart Mill (1806 – 1873) was a British philosopher and political economist. A student of Jeremy Bentham, Mill promoted the ideas of utilitarianism and liberalism and has been called “the most influential English language philosopher of the nineteenth century”. His most influential works include his books [Utilitarianism](#) (1863) and [On Liberty](#) (1859).

Longtermism

→ *Main article:* [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 utilitarianism²⁸ 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. A key reason why most utilitarians would endorse strong longtermism is that they accept *temporal impartiality*, the view that the well-being of future generations is no less important simply because they are far away in time than the well-being of those alive today. An implication of strong longtermism is to take [existential risk reduction](#) very seriously as a moral priority.

External links:

- Greaves, H. & MacAskill, W. (2019). [The case for strong longtermism](#). *Global Priorities Institute Working Paper*, 7.
- Beckstead, N. (2013). [On the Overwhelming Importance of Shaping the Far-Future](#). Ph.D. Dissertation, Rutgers University.

Maximizing utilitarianism

→ *Main article:* [Scalar versus maximizing or satisficing utilitarianism](#)

Maximizing utilitarianism is the view that within any set of options, the action that produces the most well-being is right, and all other actions are wrong.

Though this is the most common statement of utilitarianism, it may be misleading in some respects. Utilitarians agree that you *ideally* ought to choose whatever action would best promote overall well-being. That's what you have the *most* moral reason to do. But they do not recommend blaming you every time you fall short of this ideal. As a result, many utilitarians consider it misleading to take their claims about what ideally ought to be done as providing an account of moral "rightness" or "obligation" in the ordinary sense.

The main alternatives to maximizing utilitarianism include *scalar utilitarianism*, according to which rightness and wrongness are matters of degree²⁹, and *satisficing utilitarianism*, which holds that within any set of options, an action is right if it produces *enough* well-being.³⁰

Mozi

→ *Main article:* [Mozi](#)

Mò Dí (墨翟), better known as Mòzǐ or "Master Mò," flourished c. 430 BCE. in what is now Tengzhou, Shandong Province, China. Likely an artisan by craft, Mò Dí attracted many dedicated followers and founded the philosophical school of Mohism during China's Warring States Period (475 - 221 BCE)—an early predecessor to utilitarianism.

Multi-level utilitarianism

→ *Main article: [Multi-level utilitarianism versus single-level utilitarianism](#)*

Multi-level utilitarianism is the view that individuals should usually follow tried-and-tested rules of thumb, or *heuristics*, rather than trying to calculate which action will produce the most well-being. According to multi-level utilitarianism, following, under most circumstances, a set of simple moral heuristics—do not lie, steal, kill, etc.—will lead to the best outcomes overall. Often, we should use the commonsense moral norms and laws of our society as rules of thumb to guide our actions. Following these norms and laws usually leads to good outcomes because they are based on society's experience of what promotes individual well-being.

Thus, multi-level utilitarianism understands utilitarianism as a *criterion of rightness*, not as a *decision procedure*. A criterion of rightness tells us what it takes for an action (or rule, policy, etc.) to be right or wrong. A decision procedure is something that we use when thinking about what to do.

The main alternative to multi-level utilitarianism is *single-level utilitarianism*, which treats utilitarianism as both a criterion of rightness and a decision procedure.

Negative utilitarianism

Negative utilitarianism is a version of utilitarianism that assigns either no (at its most extreme) or considerably less (in its moderate form) value to the promotion of happiness relative to the reduction of suffering. One of the earliest academic formulations and critiques of negative utilitarianism was made by R. N. Smart in response to Karl Popper.³¹

External links:

- Smart, J.J.C. (1989). [Negative Utilitarianism](#), in D'Agostino F., Jarvie I.C. (eds) *Freedom and Rationality. Boston Studies in the Philosophy of Science*. 117. Springer, Dordrecht.
- Walker, A. D. M. (1974). [Negative Utilitarianism](#). *Mind*, New Series. 83(331): 424–28.
- Acton, H. B. & Watkins, J. W. N. (1963). [Symposium: Negative Utilitarianism](#). *Proceedings of the Aristotelian Society, Supplementary Volumes* 37: 83–114.

Objective list theories of well-being

→ *Main article: [Objective list theories of well-being](#)*

Objective list theories propose a list of items that constitute well-being. This list can include conscious experiences or satisfied preferences, but it rarely stops there; ethicists commonly argue that the objective list includes art, knowledge, love, friendship, and more.

The main alternatives to objective list theories include [hedonism](#), the view that well-being consists in, and only in, the balance of positive over negative conscious experiences, and [desire](#)

theories, according to which only the satisfaction of desires or preferences matters for an individual's well-being.

Objective utilitarianism

→ *Main article: [Expectation utilitarianism versus objective utilitarianism](#)*

Objective utilitarianism is the view that the rightness of an action depends on the well-being it will *in fact* produce, as opposed to the view we should promote *expected* well-being (i.e. expectational utilitarianism).

Outreach

→ *Main article: [Outreach](#)*

An effective way of doing good is by inspiring others to try to do more good. Thus, the best course of action for many people may be to develop and promote positive ideas and values, such as those associated with utilitarianism, and be a positive role-model in one's behavior. By raising awareness of positive ideas and values, it is plausible that you could inspire several people to follow their recommendations. In this way, you will achieve a multiplier effect on your social impact—the people you inspire will do several times as much good as you would have achieved by working directly to solve the most important moral problems. Because many positive ideas and values, including utilitarianism, are still little-known and little understood, there may be a lot of value in promoting them.

Peter Singer

→ *Main article: [Peter Singer](#)*

Peter Singer (1946) is an Australian moral philosopher and Professor of Bioethics at Princeton University. His work concentrates on issues in applied ethics, in particular our treatment of animals, the ethics of global poverty, and effective altruism. The publication of his 1975 book [Animal Liberation](#) helped start the modern animal rights movement.

Population ethics

→ *Main article: [Population ethics](#)*

Population ethics deals with the moral problems that arise when our actions affect who and how many people are born and at what quality of life.

Some of the main theories of population ethics include the *total view*, the *average view*, and *person-affecting views*. According to the total view, one outcome is better than another if and only if it contains greater total well-being, even if that is in virtue of simply having more people.

Similarly, according to the average view, one outcome is better than another if and only if it contains greater average well-being. Person-affecting views are a family of views that share the intuition that an act can only be good/bad if it is good/bad *for* someone. Standard person-affecting views stand in opposition to the total view since they entail that there is no moral good in bringing new people into existence because nonexistence means there is no one for whom it could be good to be created.

External links:

- Greaves, H. (2017). [Population Axiology](#). *Philosophy Compass*. 12.
- [The Repugnant Conclusion](#). The Stanford Encyclopedia of Philosophy.

Preference utilitarianism

→ *Main article:* [Theories of well-being](#)

Preference utilitarianism is the ethical theory on which one ought to promote just the sum total of preference satisfaction over dissatisfaction. In addition to the [four elements](#) shared by all utilitarian ethical theories, preference utilitarianism accepts a [desire theory of well-being](#), according to which only the satisfaction of desires or preferences matters for an individual's well-being.

Other utilitarians may accept a different theory of well-being, such as hedonism or objective list theory.

Principle of utility

In his main work *An Introduction to the Principles of Morals and Legislation*, [Jeremy Bentham](#) calls the core idea at the heart of his utilitarian philosophy the *principle of utility*. He describes it as follows: “By the ‘principle of utility’ is meant the principle that approves or disapproves of every action according to the tendency it appears to have to increase or lessen—i.e. to promote or oppose—the happiness of the person or group whose interest is in question”.³²

Prioritarianism

→ *Main article:* [Prioritarianism](#)

Prioritarianism holds that “benefiting people matters more the worse off these people are.”[*] Prioritarians thus reject the utilitarian conception of impartiality that assigns equal weight to everyone's interests (no matter their current level of well-being.)

[*]: Parfit, D. (1997). Equality and Priority. *Ratio* 10(3): 202–221, p. 213.

External links: [Priority](#), Stanford Encyclopedia of Philosophy

Quality-Adjust Life Years (QALYs)

The quality-adjusted life year (QALY) is a measure of the value of health outcomes, taking into account both *quantity* and *quality* of life.

When medical resources are scarce, utilitarians (amongst others) will want the resources to be distributed efficiently, i.e. so as to do the most good. While it would be intrusive and impractical to compare different individuals' well-being in any especially fine-grained way, it's important to at least consider the health outcomes of an intervention, such as its effects on one's life expectancy. Note that not all "life-saving" interventions are equal in this regard: to save an eighty year-old's life might really mean to provide them with 5 extra life-years (in expectation), whereas saving a thirty year-old might grant them 50+ extra life-years. This is a big difference in how much health benefit each stands to gain from having their life "saved".

But *quantity* of life is not the only thing that's relevant: we also care about *quality* of life. Health economists thus devised the *quality-adjusted life-year* metric, based on survey data of how most people would weigh trade-offs between different medical conditions and extra years of life. For example, if most people would require at least ten years of life while clinically depressed in order to outweigh the value of one year of life in full health, that suggests they value one life-year of clinical depression as roughly equal to 0.1 QALYs. If given a choice between successfully treating clinical depression for 20 years (i.e., $0.9 * 20 = 18$ QALY gain), or extending someone else's life by 10 years in full health (i.e. 10 QALY gain), these made-up numbers would suggest that the depression treatment was more important.

External links:

Sassi, F. (2006) [Calculating QALYs, comparing QALY and DALY calculations](#). *Health Policy Plan*, 21(5): 402–8. Singer, P., McKie, J., Kuhse, H., & Richardson, J. (1995). [Double jeopardy and the use of QALYs in health care allocation](#). *Journal of Medical Ethics*, 21(3): 144–150. Chappell, R.Y. (2016). [Against 'Saving Lives': Equal Concern and Differential Impact](#). *Bioethics*, 30(3): 159–164.

Richard M. Hare

→ *Main article:* [Richard M. Hare](#)

Richard M. Hare (1919 - 2002) was a British philosopher and Professor at the Universities of Oxford and Florida. One of the most influential moral philosophers of the twentieth century, Hare is most famous for his meta-ethical theory of [prescriptivism](#), which he used to argue for utilitarianism.

Rights objection to utilitarianism

→ *Main article:* [Rights objection to utilitarianism](#)

According to commonsense morality and many non-utilitarian theories, there are certain *moral constraints* you should never, or rarely, violate. These constraints are expressed in moral rules like “do not lie!” and “do not kill!”. These rules are intuitively very plausible. This presents a problem for utilitarianism. The reason for this is that utilitarianism not only specifies which outcomes are best—those having the highest overall level of well-being—but also says that it would be wrong to fail to realize these outcomes.

Sometimes, realizing the best outcome may require violating moral constraints against harming others—that is, violating their rights. For example, suppose there were five people waiting for an organ transplant and that you could save their lives if you killed one other person to harvest their organs. Intuitively, we would regard this as wrong, but it seems that utilitarianism would regard this as morally required.

See the article [The Rights Objection](#) on how proponents of utilitarianism might respond to this objection.

Rule utilitarianism ^

Rule utilitarianism is the view that what makes an action right is that it conforms to the set of rules that would have the best utilitarian consequences if they were generally accepted or followed. Since an action’s morality depends only on its conformity to a rule, rather than its own consequences, rule utilitarianism is a form of [indirect consequentialism](#).

The main alternative to rule utilitarianism is *act utilitarianism*, a direct consequentialist view, which directly assesses the moral rightness of (and only of) actions by looking at their consequences.

External links: [Rule consequentialism](#), Stanford Encyclopedia of Philosophy

Satisficing utilitarianism ^

→ *Main article:* [Scalar versus maximizing or satisficing utilitarianism](#)

Satisficing utilitarianism is the view that within any set of options, an action is right if it produces *enough* well-being.

However, this proposal has some problems and has not found wide support. To see this, suppose that Sophie could save no one, or save 999 people at great personal sacrifice, or save 1,000 people at even greater personal sacrifice. From the utilitarian’s perspective, we still want to say there is reason to save the 1,000 people over the 999 people; labeling both actions as *right* would risk ignoring the important moral difference between these two options.

The main alternatives to satisficing utilitarianism are *scalar utilitarianism*, according to which rightness and wrongness are matters of degree³³, and *maximizing utilitarianism*, the view that

within any set of options, the action that produces the most well-being is right, and all other actions are wrong.

External links:

- Bradley, B. (2006). [Against Satisficing Consequentialism](#). *Utilitas*, 18(2): 97–108.
- Chappell, R.Y. (2019). [Willpower Satisficing](#). *Noûs* 53 (2): 251–265.
- Slote, M. & Pettit, P. (1984). [Satisficing Consequentialism](#). *Proceedings of the Aristotelian Society*, Supplementary Volumes. 58: 139–163 & 165–176.

Scalar utilitarianism ^

→ *Main article:* [Scalar versus maximizing or satisficing utilitarianism](#)

Scalar utilitarianism is the view that moral evaluation is a matter of degree: the more that an act would promote the sum total of well-being, the more moral reason one has to perform that act.

³⁴ On this view, there is no fundamental, sharp distinction between ‘right’ and ‘wrong’ actions, just a continuous scale from morally better to worse.

The main alternatives to scalar utilitarianism are *maximizing utilitarianism*, the view that within any set of options, the action that produces the most well-being is right, and all other actions are wrong, and *satisficing utilitarianism*, according to which within any set of options, an action is right if it produces *enough* well-being.

External links:

- Sinhababu, N. (2018). [Scalar Consequentialism the Right Way](#). *Philosophical Studies*. 175: 3131–3144.
- Norcross, A. (2006). [The Scalar Approach to Utilitarianism](#). In West, H. (ed.), *The Blackwell Guide to Mill’s Utilitarianism*. Hoboken, New Jersey: Wiley–Blackwell. pp. 217–232.

Sentiocentrism / pathocentrism ^

→ *Main article:* [The expanding moral circle](#)

Sentiocentrism, or pathocentrism, is the view that we should extend our moral concern to all *sentient beings*, meaning every individual capable of experiencing positive or negative conscious states. Sentience is seen as the characteristic that entitles individuals to moral concern. This includes humans and probably many non-human animals, but not plants or other entities that are non-sentient.

Many consequentialist views, including utilitarianism, accept sentiocentrism. As a result, these views tend to reject [speciesism](#), the practice of giving some sentient individuals less moral consideration than others based on their species membership.

The main alternatives to sentiocentrism are *anthropocentrism*, the view that human beings deserve (overwhelmingly) greater moral concern than other beings, and *biocentrism*, which extends equal moral consideration to all living beings, including non-sentient ones like plants.

Single-level utilitarianism ^

→ *Main article: [Multi-level utilitarianism versus single-level utilitarianism](#)*

Single-level utilitarianism is the view that utilitarianism should be understood as both a criterion of rightness and a decision procedure. A criterion of rightness tells us what it takes for an action (or rule, policy, etc.) to be right or wrong. A decision procedure is something that we use when thinking about what to do.

To our knowledge, no one has ever defended single-level utilitarianism, including the classical utilitarians.³⁵ Deliberately calculating the expected consequences of all our actions is error-prone and risks falling into decision paralysis.

The main alternative to single-level utilitarianism is *multi-level utilitarianism*, the view that individuals should usually follow tried-and-tested rules of thumb, or *heuristics*, rather than trying to calculate which action will produce the most well-being. Thus, multi-level utilitarianism understands utilitarianism as a criterion of rightness, not as a decision procedure.

External links:

- Hare, R.M. (1981). Chapters 1–3, *Moral Thinking: Its Methods, Levels, and Point*. Oxford: Oxford University Press.
- Roger Crisp (1997). *Routledge Philosophy Guidebook to Mill on Utilitarianism*. Routledge, pp. 105–112.

Speciesism ^

→ *Main article: [Speciesism](#)*

Since utilitarianism accepts [impartiality](#), it considers not only the well-being of humans but also the well-being of non-human animals. Consequently, utilitarianism rejects [speciesism](#), the practice of giving individuals less moral consideration than others 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.

Consequently, 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 is 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 is 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 and deserve equal treatment.

An implication of rejecting speciesism is to take improving [farm animal welfare](#) very seriously as a moral priority.

Supererogation

→ *Main article:* [Demandingness](#)

Many ethical theories posit that some actions are *supererogatory*; that is, they are morally good but not required. In contrast, most consequentialist theories, including utilitarianism, deny that supererogatory actions exist. Utilitarianism requires us to always act such as to bring about the best outcome. The theory leaves no room for actions that are permissible yet do not bring about the best consequences. Any time you can do more to help other people than you can to help yourself, you should do so. 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. This makes utilitarianism a very [demanding](#) ethical theory.

External links: [Supererogation](#), Stanford Encyclopedia of Philosophy

Total view (population ethics)

→ *Main article:* [Total view \(population ethics\)](#)

The total view of population ethics regards one outcome as better than another if and only if it contains greater total well-being, even if that is in virtue of simply having more people.

Importantly, one population may have greater total well-being than another in virtue of having more people. One way to calculate this total is to multiply the number of individuals with their average quality of life. For example, the total view regards a world with 100 inhabitants at average well-being level 10 as just as good as another world with 200 inhabitants at well-being level 5—both worlds contain 1,000 units of well-being.

Thus, the total view implies that we can improve the world in two ways: either we improve the quality of life of existing people or we increase the number of people living positive lives. So, for example, the total view regards having a child that lives a happy and fulfilled life as something that makes the world better, other things being equal, since it adds to the total sum of well-being.³⁶ In practice, there are often trade-offs between making existing people happier and creating additional happy people. On a planet with limited resources, adding more people to an

already large population may at some point diminish the quality of life of everyone else severely enough that total well-being decreases.

The total view's foremost practical implication is [giving great importance](#) to ensuring the long-term flourishing of civilization. Since the total well-being enjoyed by all future people is potentially enormous, according to the total view, the [mitigation of existential risks](#)—which threaten to destroy this immense future value—is one of the principal moral issues facing humanity.

The main alternatives to the total view are the [average view](#), according to which one outcome is better than another if and only if it contains greater average well-being, and [person-affecting views](#), a family of views that share the intuition that an act can only be good/bad if it is good/bad *for* someone. Standard person-affecting views stand in opposition to the total view since they entail that there is no moral good in bringing new people into existence because nonexistence means there is no one for whom it could be good to be created.

External links:

Greaves, H. (2017). [Population Axiology](#). *Philosophy Compass*. 12.

[The Repugnant Conclusion](#), The Stanford Encyclopedia of Philosophy

Utilitarianism ^

→ *Main article:* [Utilitarianism](#)

Utilitarianism is the view that one morally ought to promote just the sum total of well-being.³⁷ The four elements shared by all utilitarian theories include (i) [consequentialism](#), (ii) [welfarism](#), (iii) [impartiality](#), and (iv) [aggregationism](#).

Utility ^

In philosophy, the term *utility* refers to a measure of moral value. Traditionally, utility was used to denote related concepts such as well-being, happiness, and pleasure, which are the fundamental units of value in utilitarian ethics.

In contemporary contexts, utility is predominantly used as an economic concept (as in “utility function”) to describe a person’s preference ordering over a set of alternatives.

Utility monster ^

The utility monster is a thought experiment devised by Robert Nozick to criticize utilitarianism.³⁸ Nozick imagines a hypothetical being, the utility monster, which has the capacity for generating much higher levels of well-being than anyone else. From a utilitarian perspective,

Nozick writes, the existence of such a being would require providing it with immense resources to increase its well-being, even at significant sacrifice to others.

For a utilitarian critique, see:

Chappell, R.Y. (2021). [Negative Utility Monsters](#). *Utilitas* 33 (4): 417–421.

Virtue ethics

According to virtue ethics, morality is fundamentally about having or developing a virtuous character.

The main alternatives to virtue ethics are [consequentialism](#), according to which what fundamentally matters is promoting good consequences, and [deontology](#), which views morality as being about following a system of duties and rules, like “Do Not Lie” or “Do Not Steal”.

External links: [Virtue Ethics](#), The Stanford Encyclopedia of Philosophy

Welfarism

→ *Main article:* [Welfarism](#)

Welfarism is the view that only the *welfare* (also called *well-being*) of individuals determines how good a particular state of the world is. Philosophers use the term well-being to describe everything that is good for a person in itself, as opposed to things only instrumentally good for a person. For example, money can buy many useful things and is thus good for a person instrumentally, but it is not a component of their well-being.

Welfarism is one of the [four elements of utilitarian ethical theories](#).

There are various types of welfarism, each of which regards different things as the constituents of well-being. The three most prevalent welfarist theories are [hedonism](#), [desire theories](#), and [objective list theories](#).

External links: [Welfarism](#), The Stanford Encyclopedia of Philosophy

Well-being / Welfare

→ *Main article:* [Theories of Well-Being](#)

Philosophers use the term well-being to describe everything that is good for a person in itself, as opposed to things only instrumentally good for a person. For example, money can buy many useful things and is thus good for a person instrumentally, but it is not a component of their well-being.

External links: [Well-being](#), The Stanford Encyclopedia of Philosophy

How to Cite This Page

MacAskill, W., Meissner, D., and Chappell, R.Y. (2023). Glossary. In R.Y. Chappell, D. Meissner, and W. MacAskill (eds.), <<https://www.utilitarianism.net/glossary>>, accessed 3/19/2024.

1. This definition applies to a fixed-population setting, where one's actions do not affect the number or identity of people. There are utilitarian theories that differ in how they deal with variable-population settings. This is a technical issue, relevant to the discussion of [population ethics](#). ↩
2. For a discussion of global consequentialism, see (i) Pettit, P. & Smith, M. (2000). [Global Consequentialism](#), in Brad Hooker, Elinor Mason & Dale Miller (eds.), *Morality, Rules and Consequences: A Critical Reader*. Edinburgh University Press; and (ii) Ord, T. (2009). [Beyond Action: Applying Consequentialism to Decision Making and Motivation](#). DPhil Thesis, University of Oxford. ↩
3. This definition applies to a fixed-population setting, where one's actions do not affect the number or identity of people. There are aggregationist theories that differ in how they deal with variable-population settings. This is a technical issue, relevant to the discussion of [population ethics](#). ↩
4. Harsanyi formalized his argument for utilitarianism in Harsanyi, J. (1978). [Bayesian Decision Theory and Utilitarian Ethics](#). *The American Economic Review*, 68(2), 223–228. For discussion about his proof, see Greaves, H. (2017). [A Reconsideration of the Harsanyi–Sen–Weymark Debate on Utilitarianism](#). *Utilitas*, 29(2), 175–213. ↩
5. Bostrom, N. (2003). [Astronomical Waste: The Opportunity Cost of Delayed Technological Development](#). *Utilitas*. 15(3), 308–314. ↩
6. Bostrom, N. (2003). [Astronomical Waste: The Opportunity Cost of Delayed Technological Development](#). *Utilitas*. 15(3), 308–314. ↩
7. Greaves, H. (2017). [Population axiology](#). *Philosophy Compass*. 12(11). ↩
8. Parfit, D. (1984). 143. Why We Ought to Reject the Average Principle, in [Reasons and Persons](#). Oxford: Oxford University Press. ↩
9. Arrhenius, G., Ryberg, J. & Tännsjö, T. (2017). [The Repugnant Conclusion](#). *The Stanford Encyclopedia of Philosophy*. Zalta, E. N. (ed.). ↩
10. Note that Professor William MacAskill, coauthor of this website, is a cofounder of [80,000 Hours](#). ↩

11. Ord, T. (2019). [*The Moral Imperative Towards Cost-Effectiveness in Global Health*](#), In *Effective Altruism: Philosophical Issues*. Oxford: Oxford University Press. ↩
12. Cf. MacAskill, W. (2014). [*Doing Good Better: How Effective Altruism Can Help You Make a Difference*](#). New York: Random House. Chapter 1. ↩
13. GiveWell (2019). [*Your Dollar Goes Further Overseas*](#). ↩
14. GiveWell (2019). [*Your dollar goes further overseas*](#). ↩
15. GiveWell (2019). [*Against Malaria Foundation*](#). ↩
16. Alexander, L. & Moore, M. (2020). [*Deontological Ethics*](#). The Stanford Encyclopedia of Philosophy. Zalta, E. N. (ed.). ↩
17. For a detailed philosophical discussion of effective altruism, see the 16 articles included in Greaves, H. & Pummer, T. (2019). [*Effective Altruism: Philosophical Issues*](#). Oxford: Oxford University Press. ↩
18. Ord, T. (2020). [*The Precipice: Existential Risk and the Future of Humanity*](#). London: Bloomsbury Publishing, p. 37 ↩
19. Cf. Singer, P. (1981). [*The Expanding Circle: Ethics, Evolution, and Moral Progress*](#). Princeton: Princeton University Press. ↩
20. Utilitarians of any type understand “value” in terms of well-being. ↩
21. GiveWell (2019). [*Your dollar goes further overseas*](#). ↩
22. For instance, Peter Singer’s book [*The Life You Can Save*](#) (the updated 10-year anniversary edition is available for free download) makes the case for the ethical importance of improving global health and international development. ↩
23. For a discussion of global consequentialism, see (i) Pettit, P. & Smith, M. (2000). [*Global Consequentialism*](#), in Brad Hooker, Elinor Mason & Dale Miller (eds.), *Morality, Rules and Consequences: A Critical Reader*. Edinburgh University Press; and (ii) Ord, T. (2009). [*Beyond Action: Applying Consequentialism to Decision Making and Motivation*](#). DPhil Thesis, University of Oxford. ↩
24. Bentham, J. (1789). Chapter IV: Value of a Lot of Pleasure or Pain, How to be Measured, In [*An Introduction to the Principles of Morals and Legislation*](#). ↩
25. Bentham, J. (1789). [*An Introduction to the Principles of Morals and Legislation*](#). Bennett, J. (ed.), p. 23 ↩
26. Smart, J. J. C. (1956). [*Extreme and Restricted Utilitarianism*](#). *The Philosophical Quarterly*. 6(25), p. 347. ↩

27. Bostrom, N. (2011). [Infinite Ethics](#). *Analysis and Metaphysics*. 10: 9–59. ↩
28. Cf. Greaves, H. & MacAskill, W. (2019). [The Case for Strong Longtermism](#). *Global Priorities Institute*. Section 4.1. ↩
29. More precisely: the more that an act would promote the sum total of well-being, the more moral reason one has to perform that act. ↩
30. For a discussion of this view, see Slote, M. & Pettit, P. (1984). [Satisficing Consequentialism](#). *Proceedings of the Aristotelian Society*, Supplementary Volumes. 58: 139–163 & 165–176. ↩
31. Smart, R. N. (1958). [Negative Utilitarianism](#). *Mind*. 67(268): 542–43. ↩
32. Bentham, J. (1789). [An Introduction to the Principles of Morals and Legislation](#). Bennett, J. (ed.), p. 7. ↩
33. More precisely: the more that an act would promote the sum total of well-being, the more moral reason one has to perform that act. ↩
34. Norcross, A. (2020). *Morality by Degrees: Reasons Without Demands*. Oxford University Press. ↩
35. Jeremy Bentham rejected single-level utilitarianism, writing that “it is not to be expected that this process [of calculating expected consequences] should be strictly pursued previously to every moral judgment.” Bentham, J. (1789). [An Introduction to the Principles of Morals and Legislation](#). Bennett, J. (ed.), p. 23.
- Henry Sidgwick concurs, writing that “the end that gives the criterion of rightness needn’t always be the end that we consciously aim at; and if experience shows that general happiness will be better achieved if men frequently act from motives other than pure universal philanthropy, those other motives are preferable on utilitarian principles”. Sidgwick, H. (1874). [The Methods of Ethics](#). Bennett, J. (ed.), p. 201. ↩
36. Whether or not we should have the child, however, depends also on whether this improves the total well-being more than improving the lives of existing people would, and on issues regarding resource constraints and overpopulation. ↩
37. This definition applies to a fixed-population setting, where one’s actions do not affect the number or identity of people. There are utilitarian theories that differ in how they deal with variable-population settings. This is a technical issue, relevant to the discussion of [population ethics](#). ↩
38. Nozick, R. (1974). [Anarchy, State, and Utopia](#). Basic Books. ↩