

Elements and Types of Utilitarianism

Table of Contents

1. Introduction
2. The Definition of Utilitarianism
3. The Four Elements of Utilitarianism
 - Consequentialism
 - Welfarism
 - Impartiality and the Equal Consideration of Interests
 - Aggregationism
4. The Two Elements of Classical Utilitarianism
 - Theories of Well-Being: Hedonism
 - Population Ethics: The Total View
5. Further Distinctions Among Utilitarian Theories
 - Reconstructing Rightness: Maximizing, Satisficing, and Scalar Utilitarianism
 - Expectational Utilitarianism Versus Objective Utilitarianism
 - Multi-level Utilitarianism Versus Single-level Utilitarianism
 - The Difference Between Multi-Level Utilitarianism and Rule Utilitarianism
 - Global Utilitarianism and Hybrid Utilitarianism
6. Conclusion
7. Resources and Further Reading

Introduction

As explained in [Chapter 1: Introduction to Utilitarianism](#), the core idea of utilitarianism is that we should want to improve the well-being of everyone by as much as possible. Utilitarian theories share four elements: consequentialism, welfarism, impartiality, and aggregationism. Classical utilitarianism is distinctive because it accepts two additional elements: first, hedonism as a [theory of well-being](#); second, the total view of [population ethics](#). There are several further important distinctions between utilitarian theories: we can distinguish scalar from maximizing or satisficing utilitarianism, expectational from objective utilitarianism, multi-level from single-level utilitarianism, and global from hybrid utilitarianism.

The Definition of Utilitarianism

Utilitarian theories share four defining elements:

1. Consequentialism
2. Welfarism
3. Impartiality
4. Aggregationism

Combining these, we can define utilitarianism as follows:

Utilitarianism is the view that one ought always to promote overall well-being.

Sometimes philosophers talk about “welfare” or “utility” rather than “well-being”, but they typically mean the same thing.

The Four Elements of Utilitarianism

Consequentialism

Utilitarianism is a form of [*consequentialism*](#), which we define as follows:

Consequentialism is the view that one ought always to promote better outcomes.

Consequentialism claims that you have more reason to perform an action, the better the resulting outcome would be (compared to the available alternatives). On this view, the value of the outcome is what ultimately matters, from a moral perspective. So, to evaluate whether to perform an action, we should look at its overall consequences, rather than any of its other features (such as the *type* of action that it is). For instance, when breaking a promise has bad consequences—as it usually does—consequentialists oppose it. However, breaking a promise is not considered wrong in itself. In exceptional cases, breaking a promise could be the morally best action available, such as when it’s necessary to save a life. The ends in this case justify the means.

Consequentialism’s rivals offer alternative accounts of what one morally ought to do that depend on features other than the value of the resulting outcome. For example, according to [*deontology*](#), morality is about following a system of rules, like “Do Not Lie” or “Do Not Steal”. And according to [*virtue ethics*](#), morality is fundamentally about having a virtuous character. Much of consequentialism’s appeal may stem from the conviction that *making the world a better place* is [*simply more important*](#) than any of these competing moral goals.

Even some who think that morality is fundamentally about following the right rules might appreciate a consequentialist account of *which* rules are the right ones to follow.

Direct and Indirect Consequentialism: Explaining The Difference between Act Utilitarianism and Rule Utilitarianism

When offering a consequentialist account of rightness, a common distinction in the philosophical literature is between two views called *direct consequentialism* and *indirect consequentialism*.

According to the direct view, the rightness of an action (or rule, policy, etc.) depends only on its consequences. On this view, to determine the right action from among a set of feasible options, we should directly evaluate which option has the best consequences. Act utilitarianism (or act consequentialism more generally) directly assesses the moral rightness of actions.

According to indirect consequentialism, we should instead evaluate the moral status of an action *indirectly*, based on its relationship to something else (such as a rule), which is in turn assessed in terms of its consequences. The most famous indirect view is known as *rule utilitarianism* (or [rule consequentialism](#) more generally). According to rule utilitarianism, 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 on its conformity to a rule, rather than its own consequences, rule utilitarianism is a form of indirect consequentialism.

On our definition of consequentialism, only the direct view is a genuinely consequentialist position, and rule consequentialism, despite the name, is not a type of consequentialism.¹ As Brad Hooker, a prominent rule consequentialist, argues, rule consequentialism is not plausibly motivated by a consequentialist commitment to outcomes being as good as possible: the case for rule consequentialism is instead that it impartially justifies intuitively plausible moral rules.² This marks an important difference from foundationally consequentialist theories.

Though act utilitarianism assesses only actions (rather than rules) in terms of “rightness”, it nevertheless also recognizes the importance of having strong commitments to familiar moral rules. Rules such as “don't lie” and “don't kill” are regarded as useful guidelines that [we should almost always follow in practice](#)—precisely so that our actions better achieve good outcomes and avoid harm. For further discussion and clarification, see the section on “[multi-level utilitarianism](#)”, below.

Welfarism

Consequentialists differ in what they take to be the *good* to be promoted. Utilitarians endorse [welfarism](#), which we define as follows:

Welfarism is the view that the value of an outcome is wholly determined by the well-being of the individuals in it.³

Specifically, welfarism holds that positive well-being is the only intrinsic good, and negative well-being is the only intrinsic bad. Philosophers use the term “well-being” to refer to what's good *for a*

person, as opposed to what's good *per se*, or "from the point of view of the Universe" (to use Sidgwick's poetic phrase). So, for example, one might coherently think that the punishment of an evil person is good, and moreover it's good precisely because that punishment is bad *for him*. But that would be to reject the kind of welfarism that utilitarianism accepts, on which positive well-being is always intrinsically good, and negative well-being is always intrinsically bad.

Different [theories of well-being](#) regard different things as the basic constituents of well-being. The three main theories are [hedonism](#), [desire theories](#) and [objective list theories](#). [Chapter 4](#) explains these different theories of well-being in more detail.

While every plausible view recognizes that well-being is important, some philosophers reject welfarism by claiming that other things matter in addition. For example, [egalitarians](#) may hold that inequality is intrinsically bad, even when it benefits some and harms none. Others might hold that [environmental](#) and [aesthetic value](#) must be considered in addition to well-being. Welfarists claim that these other things matter only insofar as they contribute to someone's well-being.

Impartiality and the Equal Consideration of Interests

Utilitarianism is committed to a conception of *impartiality* that builds in the *equal consideration of interests*:

Impartiality is the view that a given quantity of well-being is equally valuable no matter whose well-being it is.

As utilitarian philosopher [Henry Sidgwick](#) states: "the good of any one person is no more important from the point of view... of the universe than the good of any other".⁴ Utilitarians value the well-being of all individuals equally, regardless of their nationality, gender, where or when they live, or even their species. According to utilitarianism, *in principle* you should not even privilege the well-being of yourself or your family over the well-being of distant strangers (though there may be good *practical* reasons to do so).⁵

Not all philosophers agree that impartiality is a core feature of morality. They might hold that we are allowed, or even required, to be [partial](#) towards a particular group, such as our friends and family. Or they might advance an alternative conception of "impartiality" that does not require the equal consideration of interests. For example, [prioritarianism](#) gives extra weight to the interests of the worst-off, whoever they might be.

Aggregationism

The final common element of utilitarianism is *aggregationism*, which we define as follows:

Aggregationism is the view that the value of an outcome is given by the *sum*⁶ value of the lives it contains.⁷

When combined with welfarism and impartiality, this 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. For example, utilitarianism claims that improving five lives by some amount is five times better than improving one life by the same amount.

Some philosophers deny any form of aggregationism. They may believe, for instance, that small benefits delivered to many people cannot outweigh large benefits to a few people. To illustrate this belief, suppose you face the choice between saving a given person’s life or preventing a large group of people from experiencing mild headaches. An [anti-aggregationist](#) might hold that saving the life is more morally important than preventing the headaches, *regardless of the number of headaches prevented*. Utilitarians would reason that if there are enough people whose headaches you can prevent, then the total well-being generated by preventing the headaches is greater than the total well-being of saving the life, so you should prevent the headaches.⁸ The number of headaches we have to relieve for it to be better than saving a life might be, in practice, *extremely high*—but utilitarians, nonetheless, believe there is *some* number of headaches at which this trade-off should be made.

In practice, many individuals and policymakers appear to endorse these kinds of trade-offs. For example, allowing cars to drive fast on roads increases the number of people who die in accidents. Imposing extremely low speed limits would save lives at the cost of inconveniencing many drivers. Most people demonstrate an implicit commitment to aggregationism when they judge it worse to impose these many inconveniences for the sake of saving a few lives.

The Two Elements of Classical Utilitarianism

Above we have explained the four elements accepted by all utilitarian theories: consequentialism, welfarism, impartiality, and aggregationism. While this is useful for distinguishing utilitarian from non-utilitarian moral theories, there are also important distinctions between utilitarian theories. Depending on how a utilitarian theory is spelled out, it might have widely differing practical implications and may be more or less compelling.

The oldest and most prominent utilitarian theory is *classical utilitarianism*, which can be defined as follows:

Classical utilitarianism is the view that one ought always to promote the sum total of happiness minus suffering.

Classical utilitarianism can be distinguished from the wider utilitarian family of views because it accepts two additional elements: first, [hedonism](#), the view that well-being consists only of conscious experiences; and second, the [total view](#) of population ethics, on which one outcome is better than another if and only if it contains a greater sum total of well-being, where well-being

can be increased either by making existing people better off or by creating new people with good lives.

Theories of Well-Being: Hedonism

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

Classical utilitarianism accepts *hedonism* as a theory of well-being, which we define as follows:

Hedonism is the view that well-being consists in, and only in, the balance of positive minus negative conscious experiences.

Ethical hedonists believe that the only things good in themselves are the experiences of positive conscious states, such as enjoyment and pleasure; and the only things bad in themselves are the experiences of negative conscious states, such as misery and pain. *Happiness* and *suffering* are commonly used by philosophers as shorthand for the terms *positive conscious experience* and *negative conscious experience* respectively.

We discuss the arguments for and against hedonism—and its two major rivals, *desire theories* and *objective list theories*—in Chapter 4: [Theories of Well-Being](#).

Population Ethics: The Total View

→ Main article: [Population Ethics](#)

Classical utilitarianism accepts a population ethical theory known as the *total view*, which holds that:

One outcome is better than another if and only if it contains greater total well-being.

The total view implies that we can improve the world in two ways:⁹ either we can improve the quality of life of existing people, or we can increase the number of people living positive lives.¹⁰ 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.

Alternatives to the total view in population ethics include the *average view*, *variable value theories*, *critical level (and range) theories*, and *person-affecting views*. We explain and discuss these theories in Chapter 5: [Population Ethics](#).

Further Distinctions Among Utilitarian Theories

After selecting your preferred theory of well-being and population ethics view, you should also consider:

1. how (or whether) to construct a conception of rightness;
2. whether to focus on *actual* or *expected* consequences;
3. the role of simple heuristics, derived from utilitarianism, to guide our actions in everyday life; and
4. what forms of moral evaluation apply to rules, motives, character, and other objects of moral interest beyond actions.

Reconstructing Rightness: Maximizing, Satisficing, and Scalar Utilitarianism

Utilitarianism is most often stated in its maximizing form: 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,¹¹ 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.¹² 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.¹³

To further illustrate 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 a utilitarian perspective, the most important thing is that Sophie saves either the 999 people or the 1,000 people rather than saving no-one; the difference between Sophie's saving 999 people and 1,000 people is comparatively small. However, on the maximizing form of utilitarianism, both saving no-one and saving the 999 people would simply be labeled as "wrong". While we might well accept a maximizing account of what agents *ideally ought* to do, there are further moral claims to make in addition.

Satisficing utilitarianism instead holds that, within any set of options, an action is right if and only if it produces *enough* well-being.¹⁴ This proposal has its own problems and has not yet found wide support.¹⁵ In the case given in the previous paragraph, we still want to say there is good 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. So while we may be drawn to a satisficing account of what agents are *obliged* to do in order to meet minimal moral standards,¹⁶ this view, too, requires supplementation.

Instead, it's more popular among leading utilitarians today to endorse a form of *scalar utilitarianism*, which may be defined as follows:

Scalar utilitarianism is the view that moral evaluation is a matter of degree: the more that an act would promote well-being, the more moral reason one has to perform that act.¹⁷

On this view, there is no fundamental, sharp distinction between 'right' and 'wrong' actions, just a continuous scale from morally better to worse.¹⁸

Philosophers have traditionally conceived of maximizing, satisficing, and scalar utilitarianism as competing views. But more recently, one of the authors (Chappell) has argued that utilitarians could fruitfully accept all three, by constructing *multiple* different senses of 'should' or 'right'.¹⁹ On Chappell's "deontic pluralism" view, (i) maximizers are correct to hold that Sophie *ideally* should save all 1,000 people; (ii) satisficers may be correct to hold that saving 999 is *minimally acceptable* in a way that saving no-one is not; and (iii) scalar utilitarians are correct to hold that it's ultimately a matter of degree, and that the gain from saving 999 rather than zero dwarfs the gain from saving 1,000 rather than 999.

Expectational Utilitarianism Versus Objective Utilitarianism

Given our cognitive and epistemic limitations, we cannot foresee all the consequences of our actions. Many philosophers have held that what we ought to do depends on what we (should) believe at the time of action. The most prominent example of this kind of account is *expectational utilitarianism*.²⁰

Expectational utilitarianism is the view we should promote *expected* well-being.

Expectational utilitarianism states we should choose the actions with the highest *expected value*.²¹ 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. This approach follows [expected utility theory](#), the most widely-accepted theory of rational decision making under uncertainty. 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% x 1,000 lives) whereas the latter option saves an expected 75 lives (= 50% x 150 lives). This provides an account of *rational* choice from a moral point of view.

Objective utilitarianism, by contrast, takes the extent to which we ought to perform an action to depend on the well-being it would *in fact* produce. The contrast between the two views may be illustrated using a thought experiment:

The risky treatment: A patient has a chronic runny nose that will leave her, if untreated, with a mildly lower well-being for the rest of her life. The only treatment for her condition is very risky, with only a 1% chance of success. If successful, the treatment will cure her completely,

but otherwise it will lead to her death. Her doctor gives her the treatment, it succeeds and she is cured.

The doctor's action has—as a matter of pure chance and against overwhelming odds—led to the best outcome for the patient, and not treating the patient would have left her worse off. Thus, according to objective utilitarianism, the doctor has acted rightly. However, the action was wrong from the perspective of expectational utilitarianism. The expected value of giving the treatment, with its overwhelming odds of killing her, are much worse for the patient than not treating her at all. The doctor's decision turned out to be immensely fortunate, but it was extremely reckless and irrational given their available information.

When there is a conflict in this way between which act would be *actually* best versus which would be *expectably* best, is there a fact of the matter as to which act is “really” right? Many philosophers are drawn to the view that this is a merely verbal dispute. We can talk about the actually-best option as being “objectively right”, and the expectably-best option as “subjectively right”, and each of these concepts might have a legitimate theoretical role.²² For example, subjective rightness seems more apt to guide agents, since in real life we are often uncertain what consequences will actually result from our actions. Subjective rightness is also relevant to assessing the quality of an agent's decision-making. (We think poorly of the reckless doctor, for example.) But we should presumably *prefer* that the actually-best outcome be realized, and so will be *glad* that the doctor did as they “objectively ought”, even if they acted subjectively wrongly. Objective rightness thus tracks *what a fully informed, morally ideal spectator would want you to do*.

On this understanding, objective and expectational utilitarianism aren't truly *rival* views at all. Objective utilitarianism tells us which choices are objectively preferable, and expectational utilitarianism tells us how to make *rational* moral decisions under conditions of uncertainty.²³ Their respective claims are mutually compatible.

Multi-level Utilitarianism Versus Single-level Utilitarianism

In the literature on utilitarianism, a useful distinction is made between 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.²⁴

Utilitarians believe that their moral theory is the correct criterion of rightness (at least in the sense of what “ideally ought” to be done, as discussed above). However, they almost universally discourage using utilitarianism as a decision procedure to guide our everyday actions. This would involve deliberately trying to promote aggregate well-being by constantly calculating the expected consequences of our day-to-day actions. But it would be absurd to figure out what breakfast cereal to buy at the grocery store by thinking through all the possible consequences of buying different

cereal brands to determine which one best contributes to overall well-being. The decision is low stakes, and not worth spending a lot of time on.

The view that treats utilitarianism as both a criterion of rightness and a decision procedure is known as *single-level utilitarianism*. Its alternative is *multi-level utilitarianism*, which only takes utilitarianism to be a criterion of rightness, not a decision procedure. It is defined as follows:

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 we should, under most circumstances, follow a set of simple moral heuristics—do not lie, steal, kill etc.—expecting that this 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 overall well-being. The fact that honesty, integrity, keeping promises, and sticking to the law generally have good consequences explains why in practice [utilitarians value such things highly](#), and use them to guide their everyday actions.²⁵

In contrast, to our knowledge no one has ever defended single-level utilitarianism, including the classical utilitarians.²⁶ Deliberately calculating the expected consequences of our actions is error-prone and risks falling into decision paralysis.

Sometimes, philosophers claim that multi-level utilitarianism is incoherent. But this is not true. Consider the following metaphor provided by Walter Sinnott-Armstrong: The laws of physics govern the flight of a golf ball, but a golfer does not need to calculate physical forces while planning shots.²⁷ Similarly, multi-level utilitarians regard utilitarianism as governing the rightness of actions, but they do not need to calculate expected consequences to make decisions. To the extent that following the heuristics recommended by multi-level utilitarianism results in better outcomes, the theory succeeds.

A common objection to multi-level utilitarianism is that it is *self-effacing*. A theory is said to be (partially) self-effacing if it (sometimes) directs its adherents to follow a different theory. Multi-level utilitarianism often forbids using the utilitarian criterion when we make decisions, and instead recommends acting in accordance with non-utilitarian heuristics. However, there is nothing inconsistent about saying that your criterion of moral rightness comes apart from the decision procedure it recommends, and it does not mean that the theory fails.

The Difference Between Multi-Level Utilitarianism and Rule Utilitarianism

Multi-level utilitarianism sounds similar to the position known as rule utilitarianism, which we discussed above, and it’s easy to confuse the two. Yet, the two theories are distinct and it’s

important to understand how they differ.

Multi-level utilitarianism takes utilitarianism to be the criterion of moral rightness. This means it does not regard recommended heuristics as the ultimate ethical justification of any action, which is instead determined by the action's tendency to increase well-being. In contrast, for rule utilitarianism, conformity to a set of rules is the criterion of moral rightness: the reason an action is right or wrong is that it does or does not conform to the right set of rules.

Insofar as you share the fundamental utilitarian concern with promoting well-being, and you simply worry that deliberate pursuit of this goal would prove counterproductive, this should lead you to accept multi-level utilitarianism rather than any kind of rule utilitarianism.

Global Utilitarianism and Hybrid Utilitarianism

Most discussions of utilitarianism revolve around act utilitarianism and its criterion of right action. But we can just as well consider the tendency of other things—like motives, rules, character traits, policies and social institutions—to promote well-being. Since utilitarianism is fundamentally concerned with promoting well-being, we should not merely want to perform those *actions* that promote well-being. We should also want the motives, rules, traits, policies, institutions, and so on, that promote well-being.

This aspect of utilitarianism has sometimes been overlooked, so those who seek to highlight its applicability to things besides just actions sometimes adopt the label “global utilitarianism” to emphasize this point.²⁸

Global utilitarianism is the view that the utilitarian standards of moral evaluation apply to anything of interest.

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 morally evaluates actions. This broad focus can help the view to explain or accommodate certain supposedly “non-consequentialist” intuitions. 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.

All utilitarians should agree with this much. But there is a further question regarding whether this direct utilitarian evaluation is *exhaustive* of moral assessment, or whether there is a role for other (albeit less important) kinds of moral evaluation to be made in addition. For example, must utilitarians understand *virtue* directly as a matter of character traits that tend to promote well-being,²⁹ or could they appeal to a looser but more intuitive connection (such as representing a positive orientation towards the good)?³⁰

A challenge for *pure* global utilitarianism is that it fails to capture all of the moral evaluations that we intuitively want to be able to make. For example, imagine a world in which moral disapproval was reliably counterproductive: if you blamed someone for doing X, that would just make them stubbornly do X even more in future. Since we only want people to do more good acts, would it follow that only good acts, and not bad ones, were blame-*worthy*?

Here it's important to distinguish two claims. One is the direct utilitarian assessment that it would be good to blame people for doing good acts, and not for doing bad ones, since that would yield the best results (in the imagined scenario). But a second—distinct—claim is that only bad acts are truly blame-*worthy* in the sense of intrinsically *meriting* moral disapproval.³¹ Importantly, these two claims are compatible. We may hold *both* that gratuitous torture (for example) warrants moral disapproval, *and* that it would be a bad idea to express such disapproval if doing so would just make things worse.

This argument may lead one to endorse a form of *hybrid utilitarianism*, which we define as follows:

Hybrid utilitarianism is the view that, while one ought always to promote 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.³² For example, if someone is unwittingly anti-reliable at achieving their goals (that is, they reliably achieve the *opposite* of what they intend, without realizing it), it would clearly be *unfortunate* were they to sincerely aim at promoting the general good, and we should stop them from having this aim if we can. But their good intentions may be genuinely *virtuous* and admirable nonetheless.

Purists may object that hybrid utilitarianism is not “really” a form of utilitarianism. And, indeed, it's a hybrid view, combining utilitarian claims (about what matters and what ought to be done) with claims about virtue, praise- and blame-worthiness that go beyond direct utilitarian evaluation. But so long as these further claims do not conflict with any of the core utilitarian claims about what matters and what ought to be done, there would seem no barrier to combining both kinds of claims into a unified view. This may prove a relief for those otherwise drawn to utilitarianism, but who find pure global utilitarian claims about virtue and blameworthiness to be intuitively implausible or incomplete.

Conclusion

All ethical theories belonging to the utilitarian family share four defining characteristics: they are consequentialist, welfarist, impartial, and aggregationist. As a result, they assign supreme moral importance to promoting overall well-being.

Within this family, there are many variants of utilitarian theories. The most prominent of these is classical utilitarianism. This theory is distinguished by its acceptance of hedonism as a theory of well-being and the total view of population ethics.

There are several further distinctions between utilitarian theories: we can distinguish scalar from maximizing and satisficing utilitarianism, expectational from objective utilitarianism, multi-level from single-level utilitarianism, and hybrid from pure global utilitarianism. These distinctions can make a significant difference to how plausible one finds the resulting view, and how vulnerable it is to [various objections](#).

The next chapter discusses arguments for utilitarianism, and for consequentialism more broadly.

Next Chapter: Arguments for Utilitarianism

How to Cite This Page

MacAskill, W., Meissner, D., and Chappell, R.Y. (2023). Elements and Types of Utilitarianism. In R.Y. Chappell, D. Meissner, and W. MacAskill (eds.), <<https://www.utilitarianism.net/types-of-utilitarianism>>, accessed 3/19/2024.

Resources and Further Reading

Consequentialism

- Julia Driver (2011). *Consequentialism, New Problems of Philosophy*. José Luis Bermúdez (ed.). Abingdon: Routledge.
- Samuel Scheffler (1994). *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Oxford: Clarendon Press.
- Walter Sinnott-Armstrong (2015). *Consequentialism*. *Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.).

Welfarism & Theories of Well-Being

- Roger Crisp (2017). [Well-Being](#). *The Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.).
- Nils Holtug (2003). [Welfarism – The Very Idea](#). *Utilitas*. 15(2): 151–174.
- Shelly Kagan (1992). [The Limits of Well-being](#). *Social Philosophy & Policy*. 9(2): 169–189.

Impartiality

- Robert Goodin (1988). [What is so special about our fellow countrymen?](#) *Ethics*, 98 (4): 663–686.
- Troy Jollimore (2018). [Impartiality](#). *The Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.).

Aggregationism

- John Broome (1991). [Weighing Goods: Equality, Uncertainty, and Time](#). London: Wiley-Blackwell. Chapters 4 and 10.
- Krister Bykvist (2010). [Utilitarianism: A Guide for the Perplexed](#). London: Continuum. Chapter 5: Utilitarian Aggregation.
- Alastair Norcross (1997). [Comparing Harms: Headaches and Human Lives](#). *Philosophy & Public Affairs*. 26(2): 135–167.

Hedonism

- Roger Crisp (2006). [Hedonism Reconsidered](#). *Philosophy and Phenomenological Research*. 73(3): 619–645.
- Fred Feldman (2004). [Pleasure and the Good Life: Concerning the Nature, Varieties and Plausibility of Hedonism](#). Oxford University Press.
- Shelly Kagan (1992). [The Limits of Well-being](#). *Social Philosophy & Policy*. 9(2): 169–189.
- Ole Martin Moen (2016). [An Argument for Hedonism](#). *The Journal of Value Inquiry*. 50: 267–281.
- Andrew Moore (2019). [Hedonism](#), *The Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.).

Population Ethics

- Gustaf Arrhenius, Jesper Ryberg, & Torbjörn Tännsjö (2017). [The Repugnant Conclusion](#). *The Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.).
- Gustaf Arrhenius (2000). [Future generations: A challenge for moral theory](#). PhD Dissertation, Uppsala University.
- Hilary Greaves (2017). [Population Axiology](#). *Philosophy Compass*. 12(11).

- Johan E. Gustafsson (2022). [Our Intuitive Grasp of the Repugnant Conclusion](#). In Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns (eds.), *The Oxford Handbook of Population Ethics*. Oxford University Press.
- Michael Huemer (2008). [In Defence of Repugnance](#). *Mind*. 117(468): 899–933.
- Derek Parfit (1984). [Reasons and Persons](#). Oxford: Clarendon Press.

Maximizing, Satisficing and Scalar Utilitarianism

- Ben Bradley (2006). [Against Satisficing Consequentialism](#). *Utilitas*. 18(2): 97–108.
- Richard Y. Chappell (2020). [Deontic Pluralism and the Right Amount of Good](#). In Douglas W. Portmore (ed.), *The Oxford Handbook of Consequentialism*. Oxford University Press. pp. 498–512.
- Richard Y. Chappell (2019). [Willpower Satisficing](#). *Noûs* 53(2): 251–265.
- Alastair Norcross (2020). [Morality by Degrees: Reasons Without Demands](#). Oxford University Press.
- Alastair Norcross (2006). [The Scalar Approach to Utilitarianism](#). In Henry West (ed.), *The Blackwell Guide to Mill's Utilitarianism*. Wiley-Blackwell, pp. 217–32.
- Neil Sinhababu (2018). [Scalar consequentialism the right way](#). *Philosophical Studies*. 175: 3131–3144.

Expectational Utilitarianism Versus Objective Utilitarianism

- Roger Crisp (1997). [Routledge Philosophy Guidebook to Mill on Utilitarianism](#). Routledge, pp. 99–101.
- Peter A. Graham (2021). [Subjective Versus Objective Moral Wrongness](#). Cambridge University Press.
- Frank Jackson (1991). [Decision-theoretic consequentialism and the nearest and dearest objection](#). *Ethics*, 101(3): 461–482.

Multi-level Utilitarianism Versus Single-level Utilitarianism

- Roger Crisp (1997). [Routledge Philosophy Guidebook to Mill on Utilitarianism](#). Routledge, pp. 105–112.
- Richard M. Hare (1981). [Moral Thinking: Its Levels, Method, and Point](#). Oxford University Press.
- Peter Railton (1984). [Alienation, consequentialism, and the demands of morality](#). *Philosophy and Public Affairs*. 13(2): 134–171.

Global Utilitarianism and Hybrid Utilitarianism

- Brian McElwee (2020). [The Ambitions of Consequentialism](#). *Journal of Ethics and Social Philosophy*. 17(2).
 - Richard Y. Chappell (2012). [Fittingness: The Sole Normative Primitive](#). *Philosophical Quarterly*. 62(249): 684–704.
 - Richard Y. Chappell. [Consequentialism: Core and Expansion](#), forthcoming in D. Copp, C. Rosati, and T. Rulli (eds.). *The Oxford Handbook of Normative Ethics*. Oxford University Press.
 - Toby Ord (2009). [Beyond Action: Applying Consequentialism to Decision Making and Motivation](#). DPhil Thesis, University of Oxford.
 - Philip Pettit & Michael Smith (2000). [Global Consequentialism](#). In Hooker, B., Mason, E. & Miller, D. (eds.). *Morality, Rules and Consequences: A Critical Reader*. Edinburgh University Press.
-

1. Some other cases where labels can be misleading: herbal tea is not a type of tea; a plastic flower is not a type of flower; and the flying lemur is not a lemur and does not fly. ↩
2. For instance, Hooker writes: “The viability of this defense of rule-consequentialism against the incoherence objection may depend in part on what the argument for rule-consequentialism is supposed to be. The defense seems less viable if the argument for rule-consequentialism starts from a commitment to consequentialist assessment. For starting with such a commitment seems very close to starting from an overriding commitment to maximize the expected good. The defence against the incoherence objection seems far more secure, however, if the argument for rule-consequentialism is that this theory does better than any other moral theory at specifying an impartial justification for intuitively plausible moral rules.”

Hooker, B. (2016). [Rule Consequentialism](#), In Zalta, E. N. (ed.). *The Stanford Encyclopedia of Philosophy*. ↩

3. Assigning further value to (for instance, more equal) distributions of well-being goes beyond welfarism, by treating something other than well-being itself as intrinsically valuable. ↩
4. Sidgwick, H. (1874). [The Methods of Ethics](#). Bennett, J. (ed.), p. 186. ↩
5. See [The Special Obligations Objection](#) for further discussion. ↩
6. In principle, other aggregation methods (like multiplication or something more complex) are conceivable. But we focus here on the additive form of aggregationism, since that is by far the most common view. ↩

7. 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](#) in Chapter 5. ↩

8. Derek Parfit further argues that anti-aggregative principles implausibly endorse choices that, when iterated sufficiently many times across the population, would make everyone worse off.

Parfit, D. (2003). [Justifiability to each person](#). *Ratio*, 16(4): 368–390. ↩

9. Technically, a third possibility would be to reduce the number of negative lives. Indeed, many utilitarians support (voluntary) euthanasia, based on the recognition that a life of suffering may be worse than no life at all.

See, e.g., Singer, P. (2011). Chapter 7: Taking Life: Humans, *Practical Ethics*, 3rd ed. Cambridge University Press. ↩

10. The notion of a positive life, which is crucial for the total view, only makes sense relative to a zero point on the well-being scale. This zero point is the threshold above which life becomes “worth living”. A “neutral life”, at well-being level 0, is neither “worth living” nor “not worth living”. This may be either a life with no value or disvalue, or a life with exactly as much value as disvalue.

For discussion of the subtleties surrounding the concept of a life “worth living”, see Broome, J. (2004). [Weighing Lives](#). Oxford: Oxford University Press, pp. 66–68. ↩

11. But note that it does not straightforwardly match what Bentham and Mill meant by utilitarianism. For example, Bentham said that actions should be evaluated “according to the tendency” by which they increase or decrease well-being. Similarly, Mill argued that “actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness”.

Bentham, J. (1789). [An Introduction to the Principles of Morals and Legislation](#). Bennett, J. (ed.), p. 7.

Mill, J. S. (1863). [Utilitarianism](#). Bennett, J. (ed.), p. 5. ↩

12. Blame can be thought of as either an attitude (of moral disapproval) or an action (expressing such disapproval). Blaming actions are, for utilitarians, like any other action in that they should only be performed if they serve to promote well-being. Excessive blame would likely have bad consequences, because it discourages people from even trying. Instead, we should usually praise people who take steps in the right direction. On the moral assessment of attitudes, see the discussion of global vs hybrid utilitarianism below. ↩

13. Railton, P. (1988). [How Thinking about Character and Utilitarianism Might Lead to Rethinking the Character of Utilitarianism](#). *Midwest Studies in Philosophy* 13 (1): 398–416, p.407.
- Norcross, A. (2020). [Morality by Degrees: Reasons Without Demands](#). Oxford University Press. ↩
14. 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. ↩
15. In particular, traditional satisficing accounts struggle to offer a non-arbitrary specification of the threshold for doing “enough” good, and are vulnerable to the objection that they permit the gratuitous prevention of goodness above this threshold. Both objections are addressed in Chappell, R.Y. (2019). [Willpower Satisficing](#). *Noûs* 53(2): 251–265. ↩
16. See, e.g., Chappell, R.Y. (2020). [Deontic Pluralism and the Right Amount of Good](#). In Douglas W. Portmore (ed.), *The Oxford Handbook of Consequentialism*. Oxford University Press. pp. 498–512. ↩
17. Norcross, A. (2020). [Morality by Degrees: Reasons Without Demands](#). Oxford University Press. ↩
18. Sinhababu (2018) advances a form of scalar consequentialism that takes right and wrong to themselves be matters of degree, but the boundary point between the two is determined by conversational context rather than anything morally fundamental (or indeed genuinely significant in any way). As a result, this seems to be a merely verbal variant on the definition that we use here.
- Sinhababu, N. (2018). [Scalar Consequentialism the Right Way](#). *Philosophical Studies*. 175: 3131–3144. ↩
19. Chappell, R.Y. (2020). [Deontic Pluralism and the Right Amount of Good](#). In Douglas W. Portmore (ed.), *The Oxford Handbook of Consequentialism*. Oxford University Press. pp. 498–512. ↩
20. Jackson, F. (1991). [Decision-theoretic consequentialism and the nearest and dearest objection](#). *Ethics*, 101(3): 461–482. ↩
21. In line with the above explanation of welfarism, utilitarians of any type understand “value” in terms of well-being. ↩
22. Compare the analogous distinction between what is prudentially rational versus what is objectively best for you. Obviously, both of these normative concepts have a role to play in philosophical reflections about the perspective of self-interest. There’s no inherent barrier to using the term “ought” in relation to *either* normative role, subjective or objective, so long as one is clear about which meaning is intended. ↩

23. These are both genuine normative properties. We can ask which option has the normative property of being *what a fully informed, morally ideal spectator would want the agent to do*. And we can ask which option has the normative property of being *instrumentally rational, relative to the correct moral goals, given the evidence available to the agent*. Someone who thinks there is a real dispute between the two must think that there is some *third* normative property, of being *what you really ought to do*, that might coincide with just one of the above two properties. If we doubt that there is any such third property, then there is nothing more to dispute. We just need to be clear about which of the above two properties we mean to pick out with our “ought” talk. ↩

24. Bales, R. E. (1971). [Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?](#) *American Philosophical Quarterly*. 8(3): 257–265.

For a discussion of the multi-level view in the context of Mill’s *Utilitarianism*, see Crisp, R. (1997). [Routledge Philosophy Guidebook to Mill on Utilitarianism](#). Abingdon: Routledge, pp. 105–112. ↩

25. Hare, R. M. (1981). [Moral Thinking: Its Levels, Method, and Point](#). Oxford: Oxford University Press.

Railton, P. (1984). [Alienation, consequentialism, and the demands of morality](#). *Philosophy and Public Affairs*. 13(2): 134–171. ↩

26. 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. ↩

27. Sinnott-Armstrong, W. (2019). [Consequentialism](#), *The Stanford Encyclopedia of Philosophy*. Zalta, E. N. (ed.). ↩

28. Advocates of global consequentialism have framed it as marking a departure from traditional act consequentialism, but there has been subsequent dispute over this claim. For advocacy of global consequentialism as a distinct view, see:

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.

Ord, T. (2009). [Beyond Action: Applying Consequentialism to Decision Making and Motivation](#). DPhil Thesis, University of Oxford.

For criticism, see:

McElwee, B. (2020). [The Ambitions of Consequentialism](#). *Journal of Ethics and Social Philosophy*. 17(2).

Chappell, R. Y. (2012). [Fittingness: The Sole Normative Primitive](#). *Philosophical Quarterly*. 62(249): 684–704.

The latter argues that Global Consequentialism is best understood as a merely verbal variant of Act Consequentialism. ↩

29. Driver, J. (2001). Chapter Four: A Consequentialist Theory of Virtue. [Uneasy Virtue](#). Cambridge University Press, 63–83. ↩

30. Hurka, T. (2001). [Virtue, Vice, and Value](#). Oxford University Press. ↩

31. This mirrors Parfit's distinction between 'state-given' vs 'object-given' reasons. See Parfit, D. (2011) *On What Matters*, vol 1. Oxford University Press, p. 50. ↩

32. For further defense of this view, see Chappell, R.Y. Consequentialism: Core and Expansion, forthcoming in D. Copp, C. Rosati, and T. Rulli (eds.) *The Oxford Handbook of Normative Ethics*. ↩