

The Rights Objection

Table of Contents

1. Introduction
2. Accommodating the Intuition
3. Debunking the Intuition
4. Rivals Fare No Better
5. Biting the Bullet
6. Resources and Further Reading

Introduction

According to commonsense morality and many non-utilitarian theories, there are certain *moral constraints* you should never 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 directs us to realize these outcomes.

Sometimes, securing the best outcome may require violating moral constraints against harming others—that is, violating their rights. There is no guarantee that commonsense moral rules will always coincide with the best ways to act according to utilitarianism; we could imagine commonsense morality and good outcomes conflicting. As [previously mentioned](#), critics offer the *Transplant* thought experiment as an example of such a conflict: ¹

Transplant: Imagine a hypothetical scenario in which there are five patients, each of whom will soon die unless they receive an appropriate transplanted organ—a heart, two kidneys, a liver, and lungs. A healthy patient, Chuck, comes into the hospital for a routine check-up and the doctor finds that Chuck is a perfect match as a donor for all five patients. Should the doctor kill Chuck and use his organs to save the five others?

At first glance, it seems that utilitarianism has to answer the question with “Yes, the doctor should kill Chuck”. It’s better that five people survive than that only one does (all else equal). But on commonsense morality and virtually every other moral theory, the answer is “No, do not kill Chuck”. Killing Chuck would be widely regarded as morally monstrous. Utilitarianism seems to be

the rare exception that claims otherwise. This apparent implication is often taken to show that utilitarianism must be wrong.

Proponents of utilitarianism might [respond to this objection in four ways](#). We will go through them in turn.

Accommodating the Intuition

A first utilitarian response to the thought experiment might be to *accommodate the intuition* against killing Chuck by showing that utilitarianism does not *actually* imply that doctors should kill their patients. Critics of utilitarianism *assume* that, in Transplant, the doctor killing Chuck will cause better consequences. But this assumption is itself highly counterintuitive. If the hospital authorities and the general public learned about this incident, a major scandal would result. People would be terrified to go to the doctor. As a consequence, many more people could die, or suffer serious health problems, due to not being diagnosed or treated by their doctors. Since killing Chuck would not clearly result in the best outcome, and may even result in a terrible outcome, utilitarianism does not necessarily imply that the doctor should kill him.

Even if we *stipulate* that the scenario is an unusual situation in which killing Chuck *really would* lead to the best outcome (with no further unintended consequences), it's hard to imagine how the *doctor* could be so certain of this. Given how incredibly bad it would be to undermine public trust in our medical institutions (not to mention the reputation harm of undermining utilitarian ethics in the broader society),² it would seem unacceptably *reckless*, according to [expectational utilitarianism](#), for the doctor to risk such population-wide harm to save just a small handful of lives. Utilitarianism can certainly condemn such recklessness, even while allowing that there are rare cases in which, by unpredictable fluke, such reckless behavior could turn out to be for the best.

This is a generalizable defense of utilitarianism against a wide range of alleged counterexamples. Such “counterexamples” invite us to imagine that a typically-disastrous class of action (such as killing an innocent person) just so happens, in this special case, to produce the best outcome. But the agent in the imagined case generally has no good basis for discounting the typical risk of disaster. So it would be unacceptably risky for them to perform the typically-disastrous act.³ We maximize expected value by avoiding such risks.⁴ For all practical purposes, utilitarianism recommends that [we should refrain from rights-violating behaviors](#), just as moral intuition suggests.

Debunking the Intuition

A second strategy to deal with the Transplant case (once fleshed out with all necessary stipulations) is to debunk the intuition against killing Chuck by showing that the intuition is unreliable. It's *almost* always wrong to commit murder, and we might not be reliable at identifying the exceptions. As noted above, utilitarians tend to hold that we should cultivate strong character dispositions and

social norms against murder precisely because we are so unreliable at identifying the rare exceptions: stricter norms and dispositions will presumably lead to overall better consequences than maintaining an open attitude towards murder. But this means that our intuition against killing Chuck may just result from our having—correctly, by utilitarian lights—embraced a *useful but imperfect* moral norm against murder. While this norm is correct in the vast majority of cases, it can fail under those very exceptional circumstances in which killing someone would actually bring about the best consequences.

We may also worry that the intuition reflects [an objectionable form of status quo bias](#). However terrible it is for Chuck to die prematurely, is it not—upon reflection—equally terrible for any *one* of the five potential beneficiaries to die prematurely? Why do we find it so much easier to ignore their interests in this situation, and what could possibly justify such neglect? There are practical reasons why instituting rights *against being killed* may typically do more good than rights *to have one's life be saved*, and the utilitarian's recommended "public code" of morality may reflect this. But when we consider a specific case, there's no obvious reason why the one right should be more important (let alone five times more important) than the other, as a matter of principle. So attending more to the neglected moral claims of the five who will otherwise die may serve to weaken our initial intuition that what matters most here is just that Chuck not be killed.

Rivals Fare No Better

A third response to the Transplant case is to argue that *rival views fare no better*.

As noted above, the charge of [status quo bias](#) seems especially pressing in this context. If you asked all six people from behind [the veil of ignorance](#) whether you should kill one of them to save the other five, they'd all agree that you should. A 5/6 chance of survival is far better than 1/6, after all. And it's morally arbitrary that the one happens to have healthy organs while the other five do not. There's no moral reason to privilege this antecedent state of affairs, just because it's the status quo. Yet that's just *what it is* to grant the one a right not to be killed while refusing the five any rights to be saved. It is to arbitrarily uphold the status quo distribution of health and well-being as morally privileged, no matter that we could improve upon it (as established by the impartial mechanism of the veil of ignorance). That seems pretty bad.

Another challenge may be presented by increasing the stakes in our thought experiment:

Revised Transplant: Suppose that scientists can grow human organs in the lab, but only by performing an invasive procedure that kills the original donor. This procedure can create up to one million new organs. Like before, our doctor can kill Chuck, but this time use his body to save one million people. Should she do this?

Consider how two non-utilitarians would react to Revised Transplant. The *Moderate non-utilitarian* says that, unlike in the original case, the doctor should kill Chuck because the constraint against harming others is outweighed, since enough is at stake. The *Absolutist non-utilitarian*, on

the other hand, says that the doctor still should not kill Chuck, since no amount of benefit can outweigh the injustice of killing him.

One objection to the Moderate is that their position seems *incoherent*. The rationale underlying the prohibition against killing Chuck in Transplant should also forbid killing him in Revised Transplant. In both cases, an innocent person is sacrificed for the greater good. Another objection to the Moderate is that their position is *arbitrary*. The Moderate must draw a line past which constraint violations become permissible: for example, when the benefit is for at least one million people. But why draw the line precisely at that point, rather than higher or lower? What is so special about this particular number, 1,000,000? Yet the same question can be asked for any specific number of lives saved. The only non-arbitrary positions are that of the Absolutist, for whom there is no number of lives saved that can justify killing Chuck, and that of the utilitarian, who says that killing Chuck is justified whenever the overall benefits outweigh the costs.

Absolutism seems even more counterintuitive than utilitarianism. If we continue to increase the number of lives we could save by killing Chuck—say, from one million to one billion, and so on—it soon becomes absurd to claim that doing so is impermissible. This position appears even more absurd when we consider cases involving uncertainty. For instance, it seems the Absolutist is committed to saying it's impermissible to perform the medical procedure on Chuck, even if it had only a very small chance of killing him and is guaranteed to save millions of lives.

Biting the Bullet

The final response is to simply accept that we should—in this bizarre hypothetical situation—kill Chuck despite the intuition that killing Chuck is wrong. It's regrettable that the only way to save the five other people involves Chuck's death. Yet the right action may be to kill him since it allows the five others to continue living, any one of whom matters just as much as Chuck does. Chuck's death, while unfortunate, is *stipulated* by the thought experiment to be required to create a world where there is as much well-being as possible.

All of the standard [arguments against deontic constraints](#) become relevant at this point. For example, the [hope objection](#) flags that a benevolent observer should prefer that the five be saved, and it's hard to see how deontic moral rules could matter more than what we—or any impartial benevolent observer—should hope is done.

Of course, it's important to stress that real life outcomes can't be stipulated in advance, so in real-life cases utilitarians overwhelmingly opt to "[accommodate the intuition](#)" and reject the critic's assumption that killing innocent people leads to better outcomes.

Next: The Mere Means Objection

Other Objections to Utilitarianism

How to Cite This Page

MacAskill, W., Meissner, D., and Chappell, R.Y. (2023). The Rights Objection. In R.Y. Chappell, D. Meissner, and W. MacAskill (eds.), <<https://www.utilitarianism.net/objections-to-utilitarianism/rights>>, accessed 3/19/2024.

Resources and Further Reading

- Katarzyna de Lazari-Radek & Peter Singer (2017). [*Utilitarianism: A Very Short Introduction*](#). Oxford: Oxford University Press. Chapter 4: Objections, Section “Does utilitarianism tell us to act immorally?”.
- Krister Bykvist (2010). [*Utilitarianism: A Guide for the Perplexed*](#). London: Continuum. Chapter 8: Is Utilitarianism too Permissive?
- Shelly Kagan (1998). [*Normative Ethics*](#). Boulder, CO: Westview Press. Chapter 3.
- Shelly Kagan (1989). [*The Limits of Morality*](#). New York: Oxford University Press.
- Eduardo Rivera-López (2012). [The Moral Murderer. A \(more\) effective counterexample to consequentialism](#). *Ratio*, 25(3): 307–325.
- Judith Jarvis Thomson (1976). [Killing, Letting Die, and the Trolley Problem](#). *The Monist*. 59 (2): 204–17
- Scott Woodcock (2017). [When Will a Consequentialist Push You in Front of a Trolley?](#) *Australasian Journal of Philosophy*. 95 (2): 299–316.

1. Adapted from Thomson, J. (1976). [Killing, Letting Die, and the Trolley Problem](#). *The Monist*. 59 (2): 204–17, p. 206. ↩

2. This reputational harm is far from trivial. Each individual who is committed to (competently) [acting on utilitarianism](#) could be expected to save *many* lives. So to do things that risk deterring *many* others in society (at a population-wide level) from following utilitarian ethics is to risk *immense* harm. On the reputational costs of instrumental harm, see:

Everett, J.A.C., Faber, N.S., Savulescu, J., and Crockett, M.J. (2018). [The costs of being consequentialist: Social inference from instrumental harm and impartial beneficence](#). *Journal of Experimental Social Psychology*, 79: 200–216. ↩

3. Rivera-López, E. (2012). [The moral murderer. A \(more\) effective counterexample to consequentialism](#). *Ratio*, 25(3): 307–325.

For critical discussion, see R.Y. Chappell, [Counterexamples to Consequentialism](#). ↩

4. Even if we can somehow stipulate that the agent's first-order evidence supports believing that murder is net-positive in their case, we also need to take into account the *higher-order evidence* that most people who make such judgments are mistaken. Given the risk of miscalculation, and the far greater harms that could result from violating widely-accepted social norms, utilitarianism may well recommend that doctors adopt a *strictly* anti-murder disposition, rather than being open to committing murder whenever it *seems* to them to be for the best. ↩