



CRUCIAL CONSIDERATIONS

a blog on potentially game-changing ideas in science and philosophy



[About](#) [Science](#) [Philosophy](#) [Rationality](#) [Ethics](#)

Is Utilitarianism Too Demanding?

+12 rating, 13 votes

Posted on 13th August 2015, written by Brian Tomasik with 5 Comments



Ethics

Like 205

Tweet

Share

A common objection to utilitarianism is that the philosophy is too demanding. For instance, it might seem that we should donate all our money to those in need or devote every waking hour toward helping others. This claim is based on a misunderstanding of human willpower and decision-making. The finitude of our ability to make sacrifices for others provides all the explanation needed about why utilitarianism is not excessively demanding.

Introduction

There is an immense amount of suffering in the world:

- “About 21,000 people die every day of hunger or hunger-related causes”
- millions of people suffer from illnesses and depression
- 70 billion farm animals are raised for food each year, ⅔ in factory farms

A utilitarian goal is to reduce as much of this suffering as we can – the more the better.

This leads some observers to complain that utilitarianism is “too demanding”. For instance, it might seem that you should give away all your money to the poor, at least until you become as poor as those to whom you’re giving. Or maybe you should spend every waking hour of your life campaigning ceaselessly against cruelty to animals. These ideas appear too radical, so some moral philosophers claim utilitarianism can’t be right.

Humans have finite willpower

Imagine that you did try to work every waking hour of your life fighting poverty. Perhaps you’d even cut back on sleep so that you could have more waking hours. More hours of work imply more suffering reduced, so doesn’t utilitarianism obligate you to do this?

Here’s a plausible outcome of this scenario: Two weeks into your sleep-deprived effort, you become exhausted and fall sick. You have to stay in bed to let your body and brain recover. The next day, once you’ve regained some energy, you have a surprising negative feeling toward activism. You can’t explain why, but the thought of working more on your campaign just makes you feel irritated and depressed. You decide to take another day off for recuperation. During that day, you realize how much easier life is when you’re not pushing yourself all the time. You decide, “Screw it! Utilitarianism is too hard. I’ll adopt an easier ethical view that expects less work from me.”

In contrast, if you had taken a more moderate approach to your activism, in which you allowed yourself time for relaxation, friends, sleep, and exercise, you would have been more likely to find the process fun. You would have felt rewarded knowing you were making a difference, and you would have kept up the habit into the long term. After a few months, you would have accomplished much more than your burned-out self did.

Why utilitarianism is not excessively demanding

A common theme runs through many pieces of advice about self-exertion:

- “Slow and steady wins the race” is the moral of Aesop’s The Tortoise and the Hare fable.
- “The best is the enemy of the good” said Voltaire.
- Altruism is a marathon, not a sprint, says Robert Wiblin.
- “No one could make a greater mistake than he who did nothing because he could do only a little” is a quote attributed to Edmund Burke.

Utilitarianism recommends what will achieve the greatest reduction in suffering. Because humans are not built to make immense self-sacrifices, the greatest reduction in suffering is often attained by modest, sustainable levels of exertion.

This idea makes enough sense when seen in other contexts. Suppose you’re trying to get as many miles as possible out of your car before it breaks down. You might think you should try to drive as fast as possible because then you can get lots of miles driven within a given amount of time. But what will probably happen is that the strain of driving so fast for so long will wear out

the car parts more than if you'd driven at a modest pace. The same can be true of our bodies and minds as we apply ourselves toward a goal.

We may be tempted to think that the human will is somehow privileged, because unlike a machine, it's unbounded and limitless. This idea is a mistaken carryover from days when people believed in immaterial spirits. In fact, our minds are machines just like cars (only more complicated), and they get worn down by over-exertion.

We have several components to our motivational systems, many of which are below the level of conscious access and intentional control. The unpleasantness that we'll grow to associate with activism against poverty or animal cruelty if we do it every waking hour of our lives represents a subconscious shift in our action inclinations based on negative feedback signals. Instead, we should aim to develop positive associations with altruistic work, so that we'll be inclined to do more of it in a similar way as we're inclined to reach for an extra cookie.

If we could program a robot to act in a utilitarian fashion, we could prevent it from becoming tired or losing motivation. Humans lack this degree of control over their brain wiring. And even if such a robot existed, it would still need to expend some effort on self-maintenance, and it would still need to avoid over-exertion, just like our car does.

Practical arguments are sufficient

Few people can actually become utilitarian superheroes. Most of us will achieve the best possible results by not over-extending ourselves. But instead of taking this practical point as a sufficient resolution to the demandingness objection, some philosophers go further and aim to argue that intrinsically the scope of our duties should be limited. This may be an attempt to resolve cognitive dissonance in their minds between (a) reducing suffering seems really important but (b) I don't want to devote my life to it. It's a sort of moral rationalization.

The LessWrong community has developed a principle called *Occam's imaginary razor*, which says that when you do something you know is bad (like smoking despite it being unhealthy), you should develop a rationalization that minimizes damage to correct views of the world. For example, it would significantly disrupt your epistemological sanity if you tried to prove to yourself that smoking didn't increase risk of cancer, such as by asserting that most of the scientific literature on the topic is wrong. A much less damaging excuse would be "I don't have the motivation to quit."

We can extend Occam's imaginary razor to the moral domain and propose that if you're not going to accept a moral principle (e.g., the idea that you should devote a significant portion of your life to reducing suffering on utilitarian grounds), you should distort your moral views as little as possible in explaining why. The argument that we intrinsically lack any obligations to prevent as much suffering as we realistically can is a violation of the razor. Better would be just to say that we're selfish (like most people are), and we can only muster so much willpower to help others.

Praise and blame are instrumental

The question of whether it's morally blameworthy not to devote your whole life to reducing suffering conjures the wrong idea. Utilitarianism should not be seen as a binary morality in which you're right if you do the best possible thing and wrong otherwise. Rather, utilitarianism should be regarded more like a point counter in a video game, where you aim to accumulate as many points as you can within the bounds of reason. There's no binary "right" and "wrong". You just do the best you can.

Relatedly, the idea of a "moral obligation" is not intrinsic to utilitarianism. Talk about "duties" and "requirements" is a way humans communicate when they want to motivate others strongly to perform some action. "Rightness" and "wrongness" judgments are useful instrumentally as a way to motivate good behavior.


Thus, to call someone “morally blameworthy” unless she gives up her family and friends to devote her life to reducing suffering is a self-defeating strategy. It would be like creating a club with a \$10 million membership fee. Sure, you might get a few members, but in order to appeal to a broad audience of people who can be helping with the cause, the bar has to be much lower.

In addition, it’s a mistake to think like this: “Setting a low bar is just a way to make sure more people help, but once I joined the cause, I’d see that demanding vastly more of myself would be much better than just doing a little bit. Therefore, this cause is too demanding, and I won’t join.” This is precisely Edmund Burke’s fallacy. If imagined excessive duties prevent you from accepting utilitarianism, those excessive duties were not a utilitarian recommendation to begin with. Rather, you’re making an error.

Acknowledgements

Some of the points in this piece were inspired by Lukas Gloor, Carl Shulman, and others.

Related



Artificial Free Will
October 9, 2015
In "Ethics"



Ethics and Objectivity
November 11, 2014
In "Ethics"



Why Animals Matter I
February 5, 2015
In "Ethics"

This article has 5 comments



Alexander

Thursday 20 August 2015, 6:24 pm / Reply

Great read! Thank you! 😊



Vanitas

Thursday 20 August 2015, 10:20 pm / Reply

“Because humans are not built to make immense self-sacrifices, the greatest reduction in suffering is often attained by modest, sustainable levels of exertion.”

I don’t know. Somehow, people do make immense, long-term personal sacrifices. Aid workers, international health care professionals and social justice volunteers sacrifice huge amounts of time, energy, and financial resources—and lose tons of time that could be spent with family and friends—in order to help out. As such, your empirical claim seems somewhat undermotivated: humans

are perfectly capable of making these sacrifices. Even just financially, I could clearly give 30% of my income to charity and still be OK. Should I do so, even if it means not being able to do many things that I love to do? That's demandingness in a nutshell, and it strikes me as ad hoc to simply declare that human beings are constitutionally incapable of doing much, much more than they do.



Brian Tomasik

Monday 24 August 2015, 1:56 pm / Reply

Nice point, Vanitas. 😊

One observation is that people differ in their levels of ability to exert themselves. But more relevant to your discussion is that much of the “inability” of humans to make sacrifices is psychological and context-dependent. If thinking about making a big sacrifice would cause someone to mentally “jump ship” from altruism, then the demand in question was too big for that person at that time. Maybe later the person will be more open to making bigger sacrifices. Or if the person had lived in a different environment, she might have been more willing. But regardless of how trivial are the reasons why people “jump ship”, the fact that they do is a real constraint on what we should demand of them at a given time. As noted at the end, our moral expectations of others' behavior should be calibrated based on what achieves the best results in practice. That's the point of moral expectations.



Robert

Wednesday 9 September 2015, 12:00 pm / Reply

One important objection is that even if people could transform themselves into perfect utilitarians at the push of a button, most of us would not want to.

We want personal satisfaction, justice, wasteful beauty, sometimes even irrational indulgement – and we don't want to stop wanting them.

And I think that is evidence of the implausibility of an assumed utilitarian “correctness” in human values, or society.

This does not mean we can't add utilitarian goals into the mix. And while you are thinking as a utilitarian trying to convince people, all these other things will just look like practical obstacles. But it is still a valid objection against treating utilitarianism as meta-ethically “correct”, or to pretend it is the purpose of human society.



Brian Tomasik

Wednesday 9 September 2015, 7:11 pm / Reply

Hi Robert 😊 I agree with everything in your comment. I don't think there's such a thing as a metaethically correct moral view. This piece is merely pointing out that insofar as you do identify with making as much positive altruistic impact as possible, you shouldn't discard the project entirely only because you can't or don't want to live up to what you see as its high demands.

Leave a Reply

Enter your comment here...

A Project By

EFFECTIVE
ALTRUISM
FOUNDATION

Share And Like

Like 198 Share

Tweet

Contact

info@ea-stiftung.org
Effective Altruism Foundation
Kantstr. 71
DE-10243 Berlin, Germany

Tag Cloud

Animal Ethics Artificial Intelligence **Atheism** Cognitive
Bias Consciousness Decision Theory **Ethics Evolution**
Free Will Life Means and Ends Moral Anti-Realism Natural
Selection Philosophical Methodology Philosophy of
Language Philosophy of Mind Philosophy of
Religion Problem of Evil Psychology
Rationality Reductionism Replicators Thought
Experiments trolley problem Utilitarianism Verbal Disputes
Zombies