

Virtues for Real-World Utilitarians

Stefan Schubert & Lucius Caviola

Guest essays represent only the views of the author(s).

Table of Contents

1. Introduction
2. Moderate Altruism
3. Moral Expansiveness
4. Effectiveness-focus
5. Truth-seeking
6. Collaborativeness
7. Determination
8. Utilitarianism and Common Sense Virtues
9. Resources and Further Reading

Introduction

On the face of it, utilitarianism is extraordinarily demanding. It seems to entail extreme levels of self-sacrifice and impartiality: that we must not in any way prioritize ourselves and our family members over distant strangers. And it seems to require us to spend every waking hour obsessing over how we can maximize our positive impact on the world.

But that analysis misses that humans have psychological limitations. If utilitarians were expected to live up to such strict standards, they would risk psychological collapse. And as we will see, it would harm utilitarians' incentives, thereby reducing their output and positive impact.

Utilitarians [who are looking to apply utilitarianism in the real world](#) should therefore adopt a more sophisticated strategy. They should study the psychological obstacles to utilitarianism and prioritize overcoming those that both greatly reduce their utilitarian impact and are feasible to overcome. By contrast, they should not try to overcome obstacles that do not reduce their impact much or are not possible to overcome.

In order to overcome the most important obstacles, they should cultivate a set of utilitarian virtues. To identify these virtues, we draw on two sources. First, research on the psychology of utilitarianism and related fields of psychology. Second, lessons from the [effective altruism community](#). Though utilitarianism is distinct from effective altruism, there are many effective altruist utilitarians, and they have put a lot of thought into how to apply utilitarianism in practice. Thus, it is only natural to draw on their experiences.

Our list of utilitarian virtues is as follows.¹

- *Moderate altruism*: to partly overcome selfishness and give some resources to others.
- *Moral expansiveness*: to overcome our natural partiality against distant beneficiaries.
- *Effectiveness-focus*: to overcome our emotional numbness to the scale of a problem and our biases in favor of ineffective ways of helping others.
- *Truth-seeking*: to overcome motivated reasoning and a host of epistemic biases in order to find the most effective ways of maximizing well-being.
- *Collaborativeness*: to engage in fruitful collaboration with other utilitarians and be willing to compromise with those who hold differing

moral views.

- *Determination*: to consistently act on utilitarian principles with persistence and deliberation.

Lastly, we discuss how utilitarianism relates to common sense virtues and common sense ethics. We argue that when we move from the philosophical seminar room to the real world, utilitarianism does not say that we should harm others for the greater good in the way [it naively may seem](#). But although utilitarianism converges with common sense ethics in that regard, it departs from it in other key ways: for instance, by emphasizing the importance of caring for distant beneficiaries (moral expansiveness) and the importance of always choosing the most effective ways of helping others (effectiveness-focus).

Moderate Altruism

Utilitarianism says that everything else being equal, everyone's well-being is equally valuable. That means that we have no *intrinsic* reasons to prioritize ourselves over others. This obviously clashes strongly with our natural selfishness.

It seems both feasible and impactful to partially overcome our selfishness. To do that, utilitarians should cultivate the virtue of altruism. The more difficult question is what level of altruism they should aim for. Should they aim for extreme levels of altruism, where they, for instance, give away almost all of their money—and perhaps even one of their kidneys? Or should they settle for more moderate levels of altruism?

Extreme levels of altruism are very uncommon. Only a tiny fraction of people donate a kidney to a stranger.² Similarly, very few give away almost all of their money. That suggests that it is very difficult to be completely unselfish—that there are formidable psychological obstacles to extreme altruism.

Consequently, even though utilitarians do not have any intrinsic reasons to prioritize themselves over others, they have several instrumental reasons to do so. First, extreme levels of altruism would likely increase the risk of burnout, which would make you less productive and thus less impactful. Second, it would likely reduce the appeal of utilitarianism and turn people away from communities with many utilitarian members, such as the effective altruism community.

Third, the notion that you need to give away almost all your resources may make you less motivated to acquire new resources. It effectively functions as a 100% marginal tax rate. And like taxes, such a notion may affect your incentives. Even if you fully endorse utilitarian principles, you may not be able to suppress your selfish impulses fully over the long term. Therefore, your productivity may go down as your self-imposed “tax rate” goes up.

These considerations all support moderate altruism over extreme altruism. They are effectively saying that the extreme altruist will have fewer resources that they can give to people in need. The main counter-argument is that even if the extreme altruist has fewer resources, they will give a greater share of those resources to people in need; and that will increase their utilitarian impact.

While that is true, we need to compare this added impact with the impact that we can gain by addressing other psychological obstacles, via other virtues. As we shall see, most people can likely increase their utilitarian impact at least a hundred times through increasing the *effectiveness* of their help: for instance, by choosing to donate to the most effective charities. By contrast, increasing the *amount* that we help through transitioning from moderate to extreme levels of altruism likely makes a much smaller difference. It is also probably much less psychologically costly to help more effectively than to increase the amount of help that we give to extreme levels. For these reasons, it seems sensible to prioritize virtues that increase the effectiveness of our help—such as effectiveness-focus and truth-seeking—over extreme altruism. Thus, in our view, utilitarians should settle for moderate altruism.

Moral Expansiveness

Just like utilitarianism says that we have no intrinsic reason to prioritize ourselves over other people, so it says that we have no intrinsic reason to prioritize some beneficiaries over others when we are giving away resources. But just like extreme altruism clashes with our natural inclinations, so does this type of extreme [impartiality](#). People tend to be strongly partial in favor of their family members. Such partiality is rooted in our evolutionary history, as genes favoring those who share our genes were more likely to propagate themselves. While these preferences for our kin are not completely immutable, they have proved hard to change fundamentally. For instance, attempts at communal child-rearing at Israeli kibbutzim have largely failed, as parents wanted to retain their special, partial, relationship with their children.³

Our judgment is therefore that utilitarians need not practice such extreme impartiality, just like they need not practice extreme altruism. However, there are other forms of partiality besides that which is based on family ties. People tend to favor their compatriots over foreigners, current people over future people, humans over non-human animals, and so on. These forms of partiality seem considerably weaker, psychologically speaking, than partiality that is based on family ties. They are not rooted in shared genes, nor in personal relationships. Instead, they are rooted in membership of a group which we effectively have decided to prioritize.

One reason to believe that such partiality is more mutable is that it has weakened with time: we have, in [Peter Singer's](#) words, gradually [expanded the circle of moral concern](#) to include more groups.⁴ We have become less inclined to prioritize some people over others merely because they are members of our group. There seems to be no reason to believe that this process cannot continue.

There is also another reason why overcoming this group-based partiality seems more feasible. Impartiality with respect to your family requires you to

change your life in fundamental ways. Impartiality with respect to larger groups tends to have much more modest consequences. It does affect decisions such as what charity to donate to, what party to vote for, and potentially (depending on other factors) what job to choose. But it often leaves the private and personal side of our lives relatively unaffected. In line with that, outside of work many effective altruist utilitarians lead lives that in many ways are remarkably similar to those of most people in society.

Overcoming group-based partiality is not only feasible but also impactful. As effective altruists have shown, you can often [increase your impact by helping beneficiaries that are distant from us](#)—spatially, temporally, and biologically. For instance, our money tends to “go further overseas”: donors in rich countries can do much more good by giving to the global poor than by giving to poor people in their own countries.⁵ Likewise, because reducing the suffering of animals in factory farms is so neglected, donations to farm animal welfare charities can have an outsized impact.⁶ Lastly, some effective altruists—so-called [longtermists](#)—think that interventions that help the distant future are still much more effective.⁷

In light of these considerations, we suggest that utilitarians cultivate the virtue of *moral expansiveness*.⁸ They should continue to expand their circle of moral concern to include more distant beneficiaries. They should not discriminate in favor of closer beneficiaries, except when they have a strong personal relationship with them. According to this view, utilitarians should permit themselves to be partial in favor of people within their personal sphere.⁹ But outside of that personal sphere, they should be impartial and not discriminate against distant beneficiaries.

Effectiveness-focus

Moral expansiveness helps us overcome important obstacles to utilitarianism. But even people who are prepared to help distant beneficiaries often fail to choose the most effective ways of doing good. They are not *effectiveness-*

focused. Though effectiveness–focus may seem similar to moral expansiveness, they are psychologically distinct. In a psychological study, we found that inclination to help distant beneficiaries is at best only weakly correlated with inclination to choose the most effective ways of doing good.¹⁰

A key obstacle to effectiveness is our aversion to deprioritizing causes which feel worthy and deserving of support. Since our resources are limited, we have to make tough trade–offs in order to maximize our impact. We have to persistently prioritize the most effective opportunities to help others—even though this comes at the expense of less effective opportunities that still feel worthy of support. People are often averse to such deprioritization, which reduces the effectiveness of their help.¹¹

Another obstacle to effectiveness is that people have “pet causes”, which they prioritize even if they know that other causes are more effective. In one study, participants were informed that it is more effective to support arthritis research than to support cancer research, but most still chose to support cancer research.¹² This is a common phenomenon: people prioritize causes that they have a personal connection to or that are particularly salient and striking. For instance, our studies show that many people support disaster relief (a salient cause) even when informed that it is more effective to address recurring or permanent problems.¹³

An underlying cause of the failure to choose the most effective ways of helping others is *scope neglect*: that our feelings do not scale in proportion to the amount of suffering that we observe. Our altruistic emotions did not evolve to help us maximize impact and are not attuned to the large differences in altruistic impact between opportunities to help others in the modern world.¹⁴ Therefore, we do not necessarily get more motivated to help just because we could help more people.

Overcoming these obstacles to effectiveness is highly impactful. Studies show that the most effective charities are at least one hundred times more effective

than the typical charity.¹⁵ Thus, people who donate to the most emotionally appealing charities rather than to those that are most effective tend to have only a fraction of the impact that they could have had. The same is likely true of other types of help, such as direct work on an altruistic cause.

It also seems feasible to overcome these biases and develop the virtue of effectiveness-focus. Note that we do not suggest that people need to align their feelings with the scale of the observed suffering. It is not possible to feel a million times more for a million suffering people than for a single person. Instead, we are talking about behavior and actions. We are suggesting that you should choose how to help based on assessments of impact, rather than based on intuitive reactions. And changing behavior in this regard seems much easier than changing feelings. It is by no means trivial to be consistently focused on effectiveness. But at the same time, it seems far easier than to become, for example, fully selfless, or fully impartial in relation to one's family. In our view, many effective altruists by and large do live the virtue of effectiveness-focus, showing that it is feasible to do so.

Truth-seeking

So far, we have focused on decisions where the effectiveness of different ways of helping others is known. In such cases, moral expansiveness and effectiveness-focus can help utilitarians overcome key psychological obstacles and choose the opportunities that have the greatest impact.

But more often than not, high-impact opportunities are not lined up in front of you. Instead, you have to identify the best ways of helping others yourself. At this point, there are new obstacles to utilitarianism that must be overcome.

To identify the best ways of helping, utilitarians have to analyze all the effects of their actions and arrive at an overall estimate of their impact relative to the alternatives. It goes without saying that this is hard. This is particularly true of interventions aimed at helping distant beneficiaries—which, as we have seen,

may have the greatest utilitarian impact. For instance, assessing the effects of our current actions on the long-term future seems extraordinarily hard.¹⁶

But it is not just that the problem is intrinsically difficult. Another problem resides in our own minds. A plethora of psychological biases distract us from the truth and make it harder for us to see the world as it really is. Here we will just cover a few examples.

One of the most salient epistemic problems is our tendency to engage in *motivated reasoning*. Instead of impartially evaluating the evidence, we tend to be biased in favor of views that we find politically convenient or like for other reasons.¹⁷ We are also susceptible to *confirmation bias*, selectively seeking out evidence that supports our views while neglecting evidence that would falsify them.¹⁸ Relatedly, we tend to be *overconfident*—to overestimate our own expertise relative to that of others.¹⁹ As a result, we are often insufficiently inclined to defer to experts. For instance, donors often have little knowledge of what the most effective charities are²⁰—but instead of seeking out experts, who do know, they go with their own guesses. That obviously tends to reduce the impact of their donations.

We are also, to varying degrees, *cognitive misers*—we do not seek out evidence to the extent that we should, and we often rely on intuition when it would be more appropriate to engage in more effortful deliberative reasoning.²¹ Partly for that reason, many have not acquired the “mindware”—the concepts and the reasoning tools—that they need to estimate the relative impact of different ways of helping others. Many are unfamiliar with the concept of [expected value](#), and their grasp of probabilistic reasoning is often shallow.²²

Because of these biases and other shortcomings, utilitarians need to cultivate the virtue of truth-seeking (or “the Scout Mindset” as [Julia Galef calls it](#)).²³ They should cultivate open-mindedness, epistemic humility, and epistemic impartiality.²⁴ They should defer to experts as appropriate. They should work

hard to find and analyze relevant evidence instead of going with their gut instincts. And they should acquire the scientific mindware and thinking tools that are needed to estimate impact.

Since overcoming these obstacles would allow the utilitarian to identify higher-impact ways of helping others, it would be very impactful to do so. It also seems relatively feasible: while we cannot hope to eradicate all our epistemic biases completely, we can no doubt improve. We have already made spectacular epistemic progress over the course of history—the Scientific Revolution being a salient case in point. There is no reason to believe that progress cannot continue. The effective altruism community recognizes the importance of truth-seeking, and for that reason, they celebrate it. Since humans are social creatures, that plausibly incentivizes members to cultivate truth-seeking.

The importance of truth-seeking for utilitarianism is often underrated. Arguably it is more important for utilitarianism than for common sense morality. Unlike common sense morality, utilitarianism says that you should choose the most effective ways of helping others—and to find them, you need to be truth-seeking. In that sense, real-world utilitarianism is actually quite *epistemically demanding*. Most discussions about [utilitarianism and demandingness](#) focus on demands on our material resources, but the epistemic demands are arguably more important when you apply utilitarianism in the real world. And for most people, it may be less draining to try to improve epistemically than to give away large material resources.

Collaborativeness

Our discussion so far has highlighted psychological obstacles that prevent individual utilitarians from maximizing their personal impact. But utilitarians also need to collaborate with others to maximize their collective impact. Here, new psychological obstacles present themselves. To overcome them, utilitarians should cultivate the virtue of collaborativeness.

As Adam Smith noted, there are large economic benefits to coordination and collaboration: to trade and to specialization. We are more effective when we divide labor between ourselves, specialize on specific tasks, and trade surplus goods and services for those that we lack. That is not only true of self-interested pursuits, but also of altruistic endeavors. Utilitarians are more effective if they are part of a community. This is what motivates the many utilitarians who have chosen to become part of the effective altruism community.

The effective altruism community is intensely collaborative. They specialize and divide tasks between themselves. Some members focus on “earning to give” — to earn as much money as possible in order to fund charities. Other members work directly for those charities. Still others provide career advice to help members choose between these options. This specialization increases the community’s impact substantially.

Another advantage of forming a community is that it can make it easier to cultivate and maintain the utilitarian virtues. Within the effective altruism community, these virtues are celebrated and seen as norms. And as we have seen, people tend to be more inclined to do things if they are celebrated and supported by norms.

Utilitarians should also be open to collaborating with people with different moral views. Toby Ord, one of the founders of effective altruism, has pioneered the concept of *moral trade*.²⁵ We trade all the time to satisfy our self-interested preferences: I give you my service in return for your goods. As Ord points out, we can also trade to satisfy our moral preferences. Suppose that Alice finds veganism much more morally important than Bob, whereas Bob feels more strongly about global poverty. Then Alice and Bob can satisfy their moral preferences better if Alice promises to give to global poverty charities provided that Bob reciprocates through becoming vegan.

Utilitarians should be open to moral trade (in an extended sense) with people they disagree with. For instance, they should avoid saying and doing things that are objectionable from the point of view of other reasonable moral perspectives unless it is strictly necessary. Anecdotally, it seems that people within the effective altruism community largely do that, and that that has helped the community to do more good.

For all these reasons, you will typically have a much greater impact if you collaborate with others, and if you join an impactful community. And yet people are often much less collaborative than would be ideal. There are several psychological obstacles that impede collaboration and the formation of a community. One of them is that people have a tendency to want to pursue their own projects, according to their own wishes. Another is that people simply fail to see how greatly altruistic collaboration could increase their impact. When it comes to for-profit companies, effectiveness tends to be very salient. We can see how much profit companies make, and how much they pay their employees. The effectiveness of altruistic projects is typically much less salient. That may cause people to underestimate how effective the most effective projects are,²⁶ and thereby cause them to underestimate the impact of joining one of those projects.

There is another factor that explains why people dislike moral trade in particular. They often feel outraged or disgusted over other people's moral views and therefore refuse to cooperate or even compromise with them.²⁷ So moral trade does not come naturally to most people.

Though these obstacles are real, they do not seem insurmountable. We have a proof of concept in the effective altruism community, which does collaborate relatively well. While people will never become perfectly collaborative, effective altruists have shown that it is feasible to improve a lot on the *status quo*.

Determination

We often do not do what we on some level want to do. We plan to lose weight but keep eating too much. We plan to save for retirement but keep spending our money here and now. We endorse a stringent moral philosophy in theory but fail to act on it in practice. We suffer from *intention-behavior gaps*.²⁸

So actually [acting on utilitarian principles](#) is harder than it may seem. It is all too easy to fall back on old habits—to pursue low-impact causes or not take much altruistic action at all.

One reason is simple inertia. Another is *availability bias*—the tendency to do what is most salient and talked about.²⁹ A third is that people often take a *satisficing* attitude to doing good.³⁰ Instead of considering all the possible interventions that they could pursue, and choosing the one that is most effective (maximizing), many do a more limited search and settle for an intervention that is “good enough”.

To overcome these obstacles, utilitarians should cultivate the virtue of determination. They should actively seek out the highest-impact opportunities and avoid drifting into sub-optimal solutions. They should work as hard as necessary (though without risking burnout). And they should make sure they keep motivated over the long term, since it often takes time to have a big impact. Young people who enter a career for utilitarian reasons can typically expect their impact to peak several decades later, when they have reached more senior positions.

Without determination, there is a risk that you will not have much of an impact at all—even if you have the other utilitarian virtues. It is thus highly impactful to overcome inertia, satisficing, and related psychological obstacles by cultivating the virtue of determination. And it also seems relatively feasible. Again, it helps to have community norms. The effective altruism community celebrates determination, just like it celebrates truth-seeking: consider, for instance, the effective altruist slogan “figure out how to do the most good, *and*

then do it".³¹ This seems to help utilitarian community members to stay determined and action-oriented.

Utilitarianism and Common Sense Virtues

Much of the philosophical discussion about utilitarianism focuses on counterintuitive edge cases, where utilitarianism departs from common sense morality. In particular, there has been a lot of discussion about the trolley problem and related problems [where you can save many people by killing one](#).

³² Such edge cases are useful as tests of our intuitions about whether utilitarianism is the best or most correct moral theory. But we are interested in something else, namely how utilitarianism should be applied in the real world. And in the real world, we do not encounter such edge cases very often. Thus, they are less central to real-world utilitarianism than the philosophical discussion may make it seem.

There are real-world situations that have some similarities to the trolley problem, however: situations where we could, at first glance, do good by causing harm ("instrumental harm") or by breaking common sense norms. We could steal money to give to the poor. Or we could lie about charity effectiveness to increase donations.

While it may seem that utilitarians should engage in norm-breaking instrumental harm, a closer analysis reveals that it often carries large costs. It would lead to people taking precautions to safeguard against these kinds of harms, which would be costly for society. And it could harm utilitarians' reputation,³³ which in turn could impair their ability to do good. In light of such considerations, many utilitarians have argued that it is [better to respect common sense norms](#).³⁴ Utilitarians should adopt ordinary virtues like honesty, trustworthiness, and kindness. There is a convergence with common sense morality.

As we have seen, utilitarianism also converges considerably with common sense morality concerning altruism and impartiality. Utilitarians should not feel that they have to donate almost all of their money. And they should not force themselves to be impartial between their family and strangers.

Some generalize these insights and argue that utilitarianism converges with common sense morality more or less across the board.³⁵ This view says that although utilitarianism initially may seem like a radical departure from our pretheoretical ethical worldview, a closer analysis reveals that this is not so.

But while it is true that utilitarianism overlaps more with common sense morality than one might naively think, there is something that this argument misses. Utilitarians can massively increase their impact through cultivating some key virtues that are not sufficiently emphasized by common sense morality. Some charities are extraordinarily effective compared with the average charity, and some jobs are much higher-impact than others. To find those opportunities, utilitarians need to be unusually truth-seeking. And to take them, they need to be morally expansive, effectiveness-focused, and so on.

So we suggest that in order to be effective in the real world, utilitarians should stake out a middle way. They should by and large adopt the standard common sense virtues. But in addition to them, they should also adopt six virtues that go beyond the common sense virtues. While a utilitarian life is pretty normal in some ways, it is very different in others.

Some of our suggested virtues tend to be associated with utilitarianism. That may be especially true of moral expansiveness. But others do not have a salient link to utilitarianism, and do not tend to be associated with it. They include, in particular, truth-seeking, collaborativeness, and determination. None of these virtues are conceptually tied to utilitarianism, but empirically, it turns out that they are very important in order to maximize utilitarian impact in the real world.

A longer version of this article can be found [here](#).

About the Authors



Stefan Schubert is a researcher at the Centre for Philosophy of Natural and Social Science, London School of Economics and Political Science, who works in the intersection of moral psychology and philosophy. His research focuses on questions related to effective altruism, such as why we don't invest more in safe-guarding our common future.



Lucius Caviola is a Senior Research Fellow at the Global Priorities Institute, University of Oxford. He specializes in moral psychology and is the co-author of *Effective Altruism and the Human Mind* (Oxford University Press, 2024).

How to Cite This Page

Schubert, S., & Caviola, L. (2023). Virtues for Real-World Utilitarians. In R.Y. Chappell, D. Meissner, and W. MacAskill (eds.), *An Introduction to Utilitarianism*, <<https://www.utilitarianism.net/guest-essays/virtues-for-real-world-utilitarians>>, accessed 2/25/2025.

Resources and Further Reading

- Bazerman, Max (2020). [Better, not perfect: a realist's guide to maximum sustainable goodness](#). Harper Business.
- Bradley, Ben (2018). [Contemporary Consequentialist Theories of Virtue](#). In *The Oxford Handbook of Virtue*.
- Caviola, Lucius, David Althaus, Stefan Schubert, and Joshua Lewis (2022). [What psychological traits predict interest in effective altruism?](#) *Effective Altruism Forum*.
- Caviola, Lucius, Stefan Schubert, and Joshua D. Greene (2021). [The Psychology of \(In\) Effective Altruism](#). *Trends in Cognitive Sciences*.
- Crisp, Roger (1992). [Utilitarianism and the Life of Virtue](#). *The Philosophical Quarterly*, 42, pp. 139–160.
- de Lazari-Radek, Katarina, and Peter Singer (2013). [How much more demanding is utilitarianism than common sense morality?](#) *Revue Internationale de Philosophie*, 67(4): 427–438.
- Ord, Toby (2009). [Beyond Action: Applying consequentialism to decision making and motivation](#). D. Phil. dissertation, University of Oxford.
- Schmidt, Andreas. T. (2017). [Being Good by Doing Good: Goodness and the Evaluation of Persons](#). *Utilitas*, 29(1): 3.
- Schubert, Stefan and Lucius Caviola. [Virtues for Real-World Utilitarianism](#). In Aguiar, Fernando, Antonio Gaitán, and Hugo Viciano, *Issues in Experimental Moral Philosophy*, Routledge (forthcoming).
- Schubert, Stefan, Ben Garfinkel, and Owen Cotton-Barratt (2017). [Considering Considerateness: Why communities of do-gooders should be exceptionally considerate](#). EffectiveAltruism.org.

Videos of talks:

- Caviola, Lucius (2017). [Against Naïve Effective Altruism](#).
- Schubert, Stefan (2019). [Moral Aspirations and Psychological Limitations](#).

- Schubert, Stefan (2020). [Naïve effective altruism and the danger of neglecting psychology](#).
-

1. Importantly, this list is tentative and not meant to be exhaustive. ↩
2. Crockett, M. J., & Lockwood, P. L. (2018). Extraordinary altruism and transcending the self. *Trends in Cognitive Sciences*, 22(12): 1071–1073. ↩
3. Christakis, Nicholas A (2019). *Blueprint: The evolutionary origins of a good society*. Hachette UK. ↩
4. Singer, Peter (1981/2011). *The expanding circle: Ethics, evolution, and moral progress*. Princeton University Press. ↩
5. MacAskill, William (2015). *Doing good better: Effective altruism and a radical new way to make a difference*. Guardian Faber Publishing. ↩
6. Animal Charity Evaluators (2016). [Why Farmed Animals](#). ↩
7. Greaves, Hilary, and William MacAskill (2021). [The case for strong longtermism](#). *Global Priorities Institute, working paper*, 5. ↩
8. Crimston, Daniel, Paul G. Bain, Matthew J. Hornsey, and Brock Bastian (2016). Moral expansiveness: Examining variability in the extension of the moral world. *Journal of Personality and Social Psychology* 111(4): 636–653; Crimston, Daniel, Matthew J. Hornsey, Paul G. Bain, and Brock Bastian (2018). Toward a psychology of moral expansiveness. *Current Directions in Psychological Science*, 27(1): 14–19. ↩
9. Note that since such partiality is unique to the personal sphere, this does not mean that they should permit themselves to be nepotistic in favor of family members or friends within professional or institutional contexts. ↩

10. Caviola, Lucius, David Althaus, Stefan Schubert, and Joshua Lewis (2022). [What psychological traits predict interest in effective altruism?](#). Effective Altruism Forum. ↩
11. Tetlock, P. E. (2003). Thinking the unthinkable: Sacred values and taboo cognitions.; Caviola, Lucius, Stefan Schubert, and Jason Nemirow (2020b). The many obstacles to effective giving. *Judgment and Decision Making*, 15(2): 159–172 ↩
12. Berman, Jonathan Z., Alixandra Barasch, Emma E. Levine, and Deborah A. Small (2018). Impediments to effective altruism: The role of subjective preferences in charitable giving. *Psychological science*, 29(5): 834–844. ↩
13. Jamison, Dean et al (2006). *Disease Control Priorities in Developing Countries*, World Bank; Ord, Toby (2013). The Moral Imperative Toward Cost-Effectiveness in Global Health, Center for Global Development; Caviola, Lucius, Stefan Schubert, and Jason Nemirow (2020). The many obstacles to effective giving. *Judgment and Decision Making*, 15(2): 159–172. ↩
14. Burum, Bethany, Martin A. Nowak, and Moshe Hoffman (2020). An evolutionary explanation for ineffective altruism. *Nature Human Behaviour*, 1–13; Miller, Geoffrey (2000). *The mating mind: How sexual choice shaped the evolution of human nature*. Heinemann; Simmler, Kevin, and Robin Hanson (2017). *The elephant in the brain: Hidden motives in everyday life*. Oxford University Press; Lloyd, Elisabeth, David Sloan Wilson, and Elliott Sober (2011). Evolutionary mismatch and what to do about it: A basic tutorial. *Evolutionary Applications*, 2–4; Cosmides, Leda, and John Tooby (2006). *Evolutionary psychology, moral heuristics, and the law*. Dahlem University Press. ↩
15. Caviola, Lucius, Stefan Schubert, Elliot Teperman, David Moss, Spencer Greenberg, and Nadira S. Faber (2020). Donors vastly underestimate

- differences in charities' effectiveness. *Judgment and Decision Making*, 15(4): 509–516. ↩
16. Tarsney, Christian (2019). [The epistemic challenge to longtermism](#). Global Priorities Institute, working paper, 10. ↩
 17. Kahan, Dan M (2015). The politically motivated reasoning paradigm, part 1: What politically motivated reasoning is and how to measure it. *Emerging trends in the social and behavioral sciences: An interdisciplinary, searchable, and linkable resource*: 1–16. ↩
 18. Oswald, Margit E., and Stefan Grosjean (2004). Confirmation bias. In *Cognitive illusions: A handbook on fallacies and biases in thinking, judgment and memory*, 79. ↩
 19. Hoffrage, Ulrich (2004). Overconfidence. In *Cognitive illusions: A handbook on fallacies and biases in thinking, judgment and memory*. ↩
 20. Caviola, Lucius, Stefan Schubert, and Jason Nemirow (2020). The many obstacles to effective giving. *Judgment and Decision Making*, 15(2): 159–172. ↩
 21. Stanovich, Keith E., Richard F. West, and Maggie E. Toplak (2016). *The rationality quotient: Toward a test of rational thinking*. MIT press. ↩
 22. Pinker, Steven (2021). *Rationality: What It Is, Why It Seems Scarce, Why It Matters*. Allen Lane. ↩
 23. Galef, Julia (2021). *The Scout Mindset: Why Some People See Things Clearly and Others Don't*. Penguin. ↩
 24. Stanovich, Keith E., and Richard F. West (1997). Reasoning independently of prior belief and individual differences in actively open-minded thinking. *Journal of Educational Psychology*, 89(2): 342. ↩
 25. Ord, Toby (2015). Moral trade. *Ethics*, 126(1): 118–138. ↩

26. See Caviola, Lucius, Stefan Schubert, Elliot Teperman, David Moss, Spencer Greenberg, and N. Faber (2020). Donors vastly underestimate differences in charities' effectiveness. *Judgment and Decision Making*, 15:4: 509–516. ↩
27. Haidt, Jonathan (2012). *The righteous mind: Why good people are divided by politics and religion*. Vintage. ↩
28. Sheeran, Paschal, and Thomas L. Webb (2016). The intention–behavior gap. *Social and personality psychology compass*, 10(9): 503–518. ↩
29. Schwarz, Norbert, and Leigh Ann Vaughn (2002). The Availability Heuristic Revisited: Ease of Recall and Content of Recall as Distinct Sources of Information. Chapter in *Heuristics and Biases: The Psychology of Intuitive Judgment*, edited by Thomas Gilovich, Dale Griffin, and Daniel Kahneman, 103–19. Cambridge: Cambridge University Press. ↩
30. Misuraca, Raffaella, Palmira Faraci, Amelia Gangemi, Floriana A. Carmeci, and Silvana Miceli (2015). The Decision Making Tendency Inventory: A new measure to assess maximizing, satisficing, and minimizing. *Personality and Individual Differences*, 85: 111–116. ↩
31. Cotra, Ajeya (2017). [Introduction to Effective Altruism](#). ↩
32. Foot, Philippa (1967). The problem of abortion and the doctrine of the double effect. *Oxford Review*, 5: 5–15; Thomson, Judith Jarvis (1976). Killing, Letting Die, and the Trolley Problem, *The Monist*, 59: 204–17; Greene, Joshua D., R. Brian Sommerville, Leigh E. Nystrom, John M. Darley, and Jonathan D. Cohen (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, 293(5537): 2105–2108. ↩
33. Everett, Jim AC, Nadira S. Faber, Julian Savulescu, and Molly J. Crockett (2018). The costs of being consequentialist: Social inference from instrumental harm and impartial beneficence. *Journal of experimental social psychology*, 79: 200–216. ↩

34. Mill, John Stuart (1838). Bentham, in John M. Robson (ed.) *Collected works* (vol. X), Toronto: Toronto University Press: pp. 75–116; Sidgwick, Henry (1907/1981). *The Method of Ethics*, 7th edition. Indianapolis: Hackett; Crisp, Roger (1992). Utilitarianism and the Life of Virtue. *The Philosophical Quarterly*, 42(167): 139–160; Hooker, Brad (2002). *Ideal code, real world: A rule-consequentialist theory of morality*. Oxford University Press; Ord, Toby (2009). *Beyond Action: Applying consequentialism to decision making and motivation*. D. Phil. dissertation, University of Oxford. ↩
35. Mill, John Stuart (1861/1992). *Utilitarianism*, In *On Liberty and Utilitarianism*, Knopf: Everyman's Library, Volume 81.; Railton, Peter (1984). Alienation, Consequentialism, and the Demands of Morality. *Philosophy and Public Affairs*, 13: 134–171. ↩