

When should an effective altruist donate?

William MacAskill

Global Priorities Institute | September 2019

GPI Working Paper No. 8-2019



When should an effective altruist donate?

Introduction

Effective altruism is the use of evidence and careful reasoning to work out how to maximize positive impact on others with a given unit of resources, and the taking of action on that basis. It's a philosophy and a social movement that is gaining considerable steam in the philanthropic world. For example, GiveWell, an organization that recommends charities working in global health and development and generally follows effective altruist principles, moves over \$90 million per year to its top recommendations. Giving What We Can, which encourages individuals to pledge at least 10% of their income to the most cost-effective charities, now has over 3500 members, together pledging over \$1.5 billion of lifetime donations. Good Ventures is a foundation, founded by Dustin Moskovitz and Cari Tuna, that is committed to effective altruist principles; it has potential assets of \$11 billion, and is distributing over \$200 million each year in grants, advised by the Open Philanthropy Project.¹

Philanthropists in the effective altruism community typically donate to charities that try to improve the wellbeing of the poorest people in the world, the living conditions of livestock in factory farms, or the long-term survival and flourishing of civilisation. Highly regarded charities include the Against Malaria Foundation, which helps distribute insecticide-treated bednets, Mercy for Animals, which runs campaigns to convince large corporations to stop using eggs from caged hens, and the Blue Ribbon Study Panel on Biodefense, which works to ensure that developments in synthetic biology don't lead to catastrophic outcomes, such as a man-made pathogen causing a global pandemic.

However, for the effective altruist, a crucial issue concerns *timing*. Rather than donating as soon as she can, would she have a bigger impact if she saved her money, donating a larger amount of money at a later date? Or, if the opportunities available now are sufficiently good, should she even take out a loan so that she can donate more of her income now, and less in the future? This issue is crucial both for small individual donors in the effective altruism community — most of whom currently donate regularly each year, but some of whom save in order to donate later — and for large donors like Good Ventures, which is attempting to spend down its capital quickly, but is still saving the large majority of its wealth for a later date.²

¹ Effective altruism is not merely limited to charity, but to any way of trying to make the world a better place. In particular, led by the organisation 80,000 Hours, there is a heavy focus on career choice. There are now thousands of people around the world who have chosen their careers, at least in part, on the basis of effective altruist ideas: individuals have gone into scientific research, think tanks, party politics, social entrepreneurship, finance (in order to do good through donating), and non-profit work in order to maximize the good they can do.

² For discussion of this issue by researchers within the effective altruism community, see Julia Wise, 'Giving now vs. later: a summary', *Effective Altruism Forum*, <https://forum.effectivealtruism.org/posts/7uJcBNZhinomKtH9p/giving-now-vs-later-a-summary>, and Good

The thought that it might be most effective to invest one's money in order to donate it at a later date seems often to be regarded as a surprising or counterintuitive conclusion; this is the view of both moral philosopher Dan Moller and political philosopher Laura Valentini.³ Though many individuals might choose to save their money and, for example, donate in their will, this might be at best regarded as a compromise between altruism and self-interest, rather than warranted on the grounds of altruism alone.

However, in other contexts, we seem perfectly happy with the idea that an altruist should delay their impact. For example, it seems perfectly obvious that an 18-year old who wishes to use their life to do as much good as possible should, if they can, go to a good university in order to increase their lifetime impact. But, in doing so, they are investing four years of their time in order to have a larger impact over the rest of their life.

Similarly, the idea that foundations and universities should maintain their endowments, spending only close to the legal minimum of 5% of their endowments, seems often to be regarded as common sense. The idea that Stanford or Harvard should spend down their endowment over the next few decades would be regarded as counterintuitive. But, again, this is in effect to take the position that those donations are spent at least as well in the future as they are now.

What's more, though we seem to be intuitively unhappy with the idea that it might be most effective for an individual philanthropist to invest and donate their income at a later date, we *also* seem to find unintuitive the idea that one ought to take out a loan in order to donate even more now. Indeed, such an idea might even be regarded as fanatical. But if there are strong arguments for donating now, then why should we be so surprised if those arguments also motivate taking out a loan, at least when the interest rate is low?

Moreover, when it comes to government spending, it's well accepted that a government ought to borrow in order to be able to spend more earlier on social programs. Yet, to at least some extent, the government doesn't seem so different from a philanthropist: a significant part of the government's aim is to spend its budget in a way that will enhance the living conditions of its citizens.

Perhaps there are real differences between philanthropists, career-seekers, foundations, and governments that explain why the commonly accepted answer to the question "Give now or invest and give later?" is so different for each of them. But it's hard not to suspect that some sort of bias is at play, especially in the discrepancy between how charities spend their money and how foundations spend their money. Perhaps we simply find the status quo acceptable and find any deviations from the status quo counterintuitive. Either way, the fact that we have different attitudes to giving now versus giving later across

Ventures and Giving Now vs. Later (2016 Update), *Open Philanthropy Project Blog*, <https://www.openphilanthropy.org/blog/good-ventures-and-giving-now-vs-later-2016-update>.

³ Moller, Dan. "Should we let people starve—for now?." *Analysis* 66, no. 291 (2006): 240-247; Valentini, Laura. "On the duty to withhold global aid now to save more lives in the future." *Ethics & Global Politics* 4, no. 2 (2011).

different domains makes it particularly valuable to investigate the question of when is the socially optimal time to spend money, for it may be that we come to revisionary conclusions for one or more of those domains. Moreover, despite the importance of this question, it has received very little academic attention from either philosophers or economists.

In this chapter, I will canvass what I consider to be all the main considerations that are relevant to whether an effective altruist ought to donate now or later, explaining what each of these considerations do and do not entail. In section I, I discuss what I view as relatively minor considerations, and in section II I move on to what I view as more significant considerations. In section III I will then propose a qualitative framework to help determine, for a given cause, whether you should invest to give at a later date, whether you should give now, or whether you should take out loans in order to donate even more now.

I. Minor Considerations

Tax

For individuals, it is generally more tax-efficient to give as you earn. In the UK, for example, the Gift Aid scheme makes it easy for individuals to waive income tax on their donations: simply by ticking a box, the government will add 25% onto the donation, paying to the charity the income tax that was paid on the amount donated (if one is in the lowest tax bracket). However, the amount that the government will pay is capped at the income tax one has paid over the last four years. So if one saves until one's retirement and then donates, and the amount to be donated is sufficiently large, one will not be able to add the full 25% match from the government onto the donation.

However, this is not a strong consideration in favor of giving now, because one can easily set up a donor advised fund (DAF) that enables one to donate one's income immediately, gaining the tax benefits, but only spend the money at a later date. For this reason, in my view, individual tax considerations do not provide a strong reason in favor of giving now.

For some people, tax considerations may provide a reason for giving later, if you will predictably enter a higher tax bracket in the future. In this circumstance, one may be able to donate a greater total amount if one waits until one enters that higher tax bracket, because one would get greater tax relief at that time, which could also be donated.

Future weakness of will

Another potential argument in favor of giving now is the fact that you might not follow through with your plans to give later; perhaps when you are older you might choose to spend the money you've saved on yourself instead. If so, then you should discount any

potential future returns on investment by the chance that you will keep the money for yourself instead.⁴

However, again you have the option to donate into a DAF. Because DAFs can only pay out to registered charities, you can guarantee that you will not spend the money on yourself at a later date. (The separate issue of whether you'll change *where* you choose to donate is discussed later, in the section on changing values.)

Though weakness of the will doesn't provide an argument in favor of donating now, it might provide an argument in favor of taking out a loan in order to donate even more now. If there is a significant risk that you won't donate more in the future, then taking out a loan is a way of tying yourself to the mast, effectively turning what would be an altruistic decision for your future self (to donate to charity) into a self-interested decision (to repay the loan).

What's more neglected?

In general, because of diminishing marginal returns, one can do more to address a problem the fewer funds that have already been put towards addressing that problem.

Over time, in the USA, charitable spending has stayed approximately constant at about 2% of GDP.⁵ However, average income is rising over time as the world gets wealthier, so the amount of total charitable donations is increasing; in the US it has risen from \$75 billion in inflation-adjusted dollars in 1975 to \$410 billion in 2017.⁶ This suggests that in general, the sorts of problems one would address with charity will receive more money in the future compared to what they receive now. This gives a mild consideration in favor of giving now.

Uncertainty

⁴ The relevance of weakness of the will to the question of giving now vs later is actually dependent on a contested normative assumption. According to *actualism*, you should take future weakness of the will into account in your decisions. On actualism, weakness of the will is therefore relevant to the question of giving now vs later. According to *possibilism*, however, you should not take future weakness of the will into account. Possibilism endorses (whereas actualism rejects) the rule of deontic logic on which Ought(A & B) entails Ought(A) and Ought(B). Suppose it is true that you ought to both invest your money, and donate those savings at a later date, but that, as a matter of fact, you won't donate. According to possibilism, Ought(Save now & donate later) entails Ought(Save now), and the fact that you won't actually donate later is irrelevant. I don't dwell on this issue in part for focus and in part because I find actualism to be the far more plausible position.

⁵ Havens, John J., Mary A. O'Herlihy, and Paul G. Schervish. "Charitable giving: How much, by whom, to what, and how." *The nonprofit sector: A research handbook 2* (2006): 542-567; Charity Navigator, 'Giving Statistics', <https://www.charitynavigator.org/index.cfm?bay=content.view&cpid=42>.

⁶ 'Giving USA 2018: The Annual Report on Philanthropy for the Year 2017,' Giving USA Foundation 2018.

We don't know for certain what the future will be like. Every year, it's possible that some catastrophe, such as an asteroid strike, a world war or huge depression, will end the human race, destroy civilization, or otherwise make one's current income close to worthless. This ongoing risk gives a mild consideration in favor of giving now. The Stern Review on the Economics of Climate Change, for example, incorporates a 0.1% annual discount factor into their cost-benefit analysis as a modelling assumption that represents the possibility of human extinction.⁷

Self-interest

Saving money provides benefits to the saver, effectively providing insurance in the case of major negative life events (such as unexpected illness or job loss) that require emergency uses of money. Insofar as very few people give as much as would maximize their altruistic impact, this gives a mild reason in favor of giving later.

For example, consider someone who faces the following two options:

- A. Donate yearly, giving a total of \$100,000 (inflation adjusted, i.e. in US 2019 dollars) over their lifetime.
- B. Save during their lifetime, and donate \$110,000 (inflation adjusted, i.e. in US 2019 dollars) at the end of their life unless a one-in-a-hundred negative life event occurs, in which case they spend the money on themselves.

It may be that the decision-maker prefers B to A on self-interested grounds — because she's happy to donate an extra \$10,000, with 99% probability, in order to be able to not-donate in the 1% likely outcome of a major negative life event — and that B is better than A on altruistic grounds, because B results in an expected \$108,900 (i.e. 99% * \$110,000) going to charity, whereas A results in an expected \$100,000 going to charity. So some 'give later' options could be better both in terms of self-interest and altruism than 'give now' options.

However, donating immediately has benefits, too. By being able to see the impact of your donation, you get a 'warm glow' that studies suggest improves happiness.⁸ The personal benefits of donating give a reason for spreading out your donation over the course of your life, rather than waiting until the end of your life. For this reason, I think that self-interest

⁷ Stern, Nicholas Herbert. *The economics of climate change: the Stern review*. Cambridge University Press, 2007.

⁸ Andreoni, James. "Impure altruism and donations to public goods: A theory of warm-glow giving." *The economic journal* 100, no. 401 (1990): 464-477; Dunn, Elizabeth W., Lara B. Aknin, and Michael I. Norton. "Spending money on others promotes happiness." *Science* 319, no. 5870 (2008): 1687-1688.

provides at most only a mild consideration either way (though the extent of this consideration will necessarily be person-dependent).

II. Major considerations

Valuing the present generation more highly

One reason you might have for preferring to donate now is that you value the interests of people alive today over the interests of people in the future.⁹

Moral philosophers typically argue that mere location in time provides no moral reason to discount the interests of future people. If I placed a bomb in some woods and primed it to go off in one hundred years' time, knowing that it will kill a child who is not yet born, that action seems no less wrong than if I were to prime the bomb to explode and kill a child in one week's time.

However, there might be relevant moral considerations that correlate with temporal location, even if they are, morally speaking, independent of it. For example, many people believe that they have a greater duty to their co-nationals than they have to strangers, perhaps because they are in a reciprocal relationship with co-nationals, receiving benefits from the government that are paid from taxes of co-nationals. If this gives a reason for preferring to help one's co-nationals than to help people in other countries, then it would also provide a reason for helping co-nationals that are alive today than co-nationals in the future.

One might think that the idea that some people's interests count for more than others is absurd, and therefore want to reject this consideration out of hand. But that would be too hasty. In other work,¹⁰ I have defended the idea that, given the uncertainty we face about which moral view is correct, rather than simply following the moral view that we think is most likely to be right we should instead look at all the moral views that we have some degree of confidence in, and take the action that represents the best 'hedge' between those views. More precisely, I argue that we should *maximize expected choice-worthiness*, where the expected choice-worthiness of an option is given by the sum, over all possible moral views, of the product of one's credence in each moral view and the choice-worthiness of that option if that moral view were true.

⁹ Note that if one were to act on this consideration, one would be departing from pure effective altruism, which treats all individuals' interests equally.

¹⁰ MacAskill, William, and Toby Ord, "Why maximize expected choice-worthiness?", *Nous*, 2018; MacAskill, William, "Practical Ethics Given Moral Uncertainty," *Utilitas*, 2019.

If this is true, then any reasonable person should place *some* extra weight on the interests of those to which one has a special relationship. This is because you should have some degree of confidence that all people's interests should be treated equally, and some degree of confidence that you should weight the interests of those to whom one has special relationships more heavily; both are popular views. However, you should have almost no credence in the idea that you should weight the interests of strangers *more* heavily than you weight the interests of friends, family or co-nationals. In which case, if you take an expectation over all moral theories in which you have some degree of confidence, then you will give some additional weight to the interests of people to whom you have a special relationships; you should give them a weight in between that given to them by the 'special obligations' view and the 'impartiality' view. Exactly how much weight depends on the level of credence that one gives to the view that special relationships do matter morally. Personally, I have a fairly high degree of belief that when it comes to philanthropy we should be impartial across all creatures, so this consideration does not affect my decisions too greatly. But others may have different views.

Changing opportunities

The opportunities that you will have for making an impact will change over time. If the cost-effectiveness of the best opportunities available to you is decreasing over time, then that provides an argument for giving now (or taking out a loan to give more now). If the cost-effectiveness of the best opportunities available to you is increasing over time, then that provides an argument for giving later.

The world is getting richer. Over the last 25 years, PPP-adjusted per-capita GDP has increased from \$5410 to \$15,500 — a rate of 4.3% per year.¹¹ In that same time, the number of people living in extreme poverty has decreased from 1.9 billion to only 1 billion, even though in that time the total world population increased by 35%.¹²

The fact that the world is getting richer means that in general the best giving opportunities will be drying up, because it means that those who are potential recipients of philanthropy are less badly off. They will have solved the most pressing of their own problems. And, in general, the better off someone is, the more resources it takes to give that person a given amount of benefit.

However, it is not true for every problem that a philanthropist might want to tackle that we're making progress. The real income of the poorest 20% of families in the US has barely increased since 1970;¹³ nor has real GDP per capita increased since 1980 in some

¹¹ <http://data.worldbank.org/indicator/NY.GDP.PCAP.PP.CD?end=2015&start=1961&view=chart>; data is from 1990 to 2015.

¹² <https://ourworldindata.org/world-poverty/#poverty-by-world-regions-in-absolute-numbers>; data is from 1990 to 2011.

¹³ <http://www.cbpp.org/research/poverty-and-inequality/a-guide-to-statistics-on-historical-trends-in-income-inequality>

of the world's poorest countries, such as Central Africa Republic, Liberia or Niger.¹⁴ Moreover, some global problems are a result of our increasing prosperity. The amount of ecological destruction and species loss increases as population grows and as more greenhouse gases are emitted. The number of animals raised for food increases every year: in the US the number has increased from 1.8 billion animals in 1960 to 9.2 billion in 2015.¹⁵ In these cases, we should expect to see less of an effect of the best opportunities decreasing over time.

A second way in which the cost-effectiveness of the best opportunities might change over time is based on the amount of resources going into a particular cause-area. If the amount of resources dedicated to tackling a problem is increasing, then, other things being equal, the cost-effectiveness of the best interventions available will decrease; if the amount of resources dedicated to tackling a problem is decreasing, then, other things being equal, the cost-effectiveness of the best interventions available will increase. This is because, when the amount of money flowing into tackling a problem increases, that spending will use up the room for more funding of the most effective interventions. For example, smallpox was the first disease to be eradicated; it was also significantly easier to eradicate than other diseases, because every infected person shows signs of infection, those signs are very distinctive, the initial appearance of symptoms occurs relatively soon after infection, there is only a human reservoir, and the smallpox vaccine was highly effective and did not require refrigeration. This exceptionality of smallpox is even more glaring when we consider the cost-effectiveness of the eradication effort: it was comparatively easy to eradicate, yet it was the disease that has killed the most people ever. It cost only \$1.6 billion in today's money, saving between 60 and 120 million lives—a cost-effectiveness of \$27 per life saved. In contrast, efforts to eradicate other diseases such as polio and measles have proven to be more difficult and more expensive.

This consideration is even stronger when we consider increases in resources that are *value-aligned*: that is, where the donors in question do not just support your preferred cause, but also share your values and general worldview. In this case, if you conclude that a particular cause or intervention isn't that promising, or if you discover a new cause or intervention that is particularly promising, then other donors who share your values and worldview are reasonably likely to support that too. When there is other donor money that is value-aligned, you should think of your spending as acting in harmony with the rest of the value-aligned money. If you think that the amount of value-aligned money is going to significantly increase over time, that gives an argument from diminishing marginal returns in favor of giving more now. If you think that the amount of value-aligned money is going to significantly decrease over time, that gives some reason for donating later.

A second line of thought relating to how the cost-effectiveness of the best opportunities might change over time is that entirely new opportunities might arise. For example, gene drive technology has recently been developed, potentially giving us the opportunity to

¹⁴ <http://data.worldbank.org/indicator/NY.GDP.PCAP.CD>

¹⁵ http://www.humanesociety.org/news/resources/research/stats_slaughter_totals.html

completely eradicate malaria for what in global terms would be a very small amount of money. Of course, whether the use of gene drives to eradicate malaria is feasible and safe is still an open question. But suppose that it were discovered to be safe and feasible. In that case, donors would have an entirely new funding opportunity; if it were the case that the funding opportunity were open, it would be possible that those who are interested in donating to prevent deaths from malaria would have been able to do more good by saving their money to donate to the gene drive-based eradication effort than by donating to a charity like Against Malaria Foundation that distributes bednets.

Similarly, new organizations might come into existence. Perhaps there is some program that is known to be more effective than any other, but no existing charity that implements that program. In which case, it might be best to save in order that, if and when a charity comes into existence that implements that program, one can fund it.

So we have two effects: the world is getting richer, which decreases the cost-effectiveness of the best interventions; but we are making technological and organizational progress, which generates new opportunities for doing good. It's difficult to know which effect is more important, but my guess is that it's the former. When we look at the history of philanthropy's successes, many occurred decades ago, when the lowest-hanging fruit had not been picked: examples include Rockefeller's funding for short-stem disease-resistant wheat that precipitated the Green Revolution; Katherine McCormick's funding of the contraceptive pill; and the funding behind smallpox eradication.

One possible strategy, however, is what I'll call "watch then pounce". On this strategy, you invest your money rather than spend it, watching to see if any new unusually cost-effective opportunities arise. When an unusually cost-effective opportunity does arise, you then completely fill that organization's room for more funding. This is a strategy that could be optimal even if, on average, the best opportunities were decreasing in cost-effectiveness over time.

For example, consider a very simple model where the starting cost-effectiveness of the best donation opportunity is 1. Every year you think it's 10% likely that some new opportunity will arise. If the new opportunity does arise, you will be able to fund that opportunity, which has a cost-effectiveness of 2. If the new opportunity doesn't arise, then the cost-effectiveness of the best donation opportunity will decrease by 20%.

Now suppose that you choose the 'watch then pounce' strategy. You save your money on the possibility that the new opportunity arises, then fund that opportunity when it does. After 7 years, there's a cumulative 52% probability of a new opportunity arising, so the expected value of waiting for 7 years and funding the new opportunity if and when it comes along is $0.52 * 2 = 1.04$, which is greater than the expected value of simply donating immediately (which equals 1). The expected value of 'watch then pounce' increases indefinitely, so if you can wait for 7 years or more, the 'watch then pounce' strategy does better than donating immediately. This is so even though the *expected* cost-effectiveness of the best donation opportunity is always decreasing.

It can therefore make sense to invest and give later even if cost-effectiveness is decreasing over time, depending on how quickly one thinks that cost-effectiveness over time is decreasing, and how much variance there is in the downward trajectory of cost-effectiveness over time. If cost-effectiveness is only decreasing slowly, or if variance is high, then it can make sense to wait until some new outstanding giving opportunity arises.

Getting better knowledge

Another consideration is how one's state of information about different giving opportunities will change over time. Whereas the previous consideration related to how good one's opportunities are, this consideration relates to how good one is at identifying good giving opportunities. All other things being equal, if you have more knowledge of the world, then you'll be able to do better at finding the best giving opportunities.

Our ability to identify the best giving opportunities has certainly increased over time. Over the past two decades, there has been a rise in the use of randomized controlled trials, social science in general has progressed, and new rigorous charity evaluators such as GiveWell have been founded. For almost everyone, this progress happens independently of one's donations; you get the new information almost for free. All other things being equal, therefore, it would make sense to wait as one's ability to pick the most cost-effective interventions improved in light of one's improved epistemic state.

It's unlikely that one's state of information will get worse over time. So in general this consideration should only push us in the direction of investing to give later. However, sometimes it might be the case that one can *gain* new information by giving now. You can do this by funding research, or by funding a fledgling non-profit to give it the capacity to provide you with more information, or as a way of gaining access to organizations that you might not otherwise have access to.

Values changes

A similar consideration to one's state of knowledge is how one's values might change over time.

It's an interesting philosophical issue to what extent we should trust the values of our future selves. Most of us would agree that you should trust the views of your future self over your own when it comes to empirical beliefs, at least absent brain injuries or any other mitigating conditions, and on the assumption that your future self has had more time to investigate the issue than you have. But when it comes to values, people tend to be more uneasy. It's a common trope that people become more conservative over time, and suppose that you knew that in the future you would be significantly more conservative than you are today. Should this affect your decision about whether to give

now or later? On the one hand, you might think that your future self has had more time to reflect and listen to new moral arguments, and therefore is in a better position to come to a set of values than you are today; on this view, the fact that your values might change over time provides a reason for you to invest, so that you can spend the money in a way that's more in line with the correct moral view.

On the other hand, you might see this as 'value-drift', providing a reason for you to donate sooner; on this view, it's your *current* values that you should try to optimize for, even if your future self has had more time to reflect. On this view, you should give now, or even take out a loan in order to prevent your future self from spending their money in a poor way.

My own view is that we should defer to our future selves on our values, *if* it's the case that we think that their values are the result of 'good' processes, such as careful reflection, discussion with peers, consideration of moral argument, and so on. There are, however, many value-changes that are the result of bias, or cognitive decline, or personality changes, rather than of genuine reflection. For example, if as your income increases you were to come to the view that there are no strong altruistic obligations, then you might justifiably suspect that your future self's moral views have been biased by self-interest.

Investment returns

The final consideration we'll discuss is the rate of return that you can get by investing the money. The real rate of return of the S&P over the last hundred years has been about 7%; this is almost certainly higher than the expected return one can get going forward (because the S&P did unusually well in this time period, and there may be some very small annual chance of catastrophic losses that are not captured by one hundred years of data). However, the rate of return on investment is always positive on average, and so by investing the money in an index fund or other vehicle, you can therefore ultimately donate a much larger amount of money than you started with. This may seem like a simple consideration in favor of giving later. But the reality is actually fairly complicated.

It's been argued that, just as one gets a financial return from investing one's money, one gets a social return by donating;¹⁶ making the extreme poor richer, for example, has compounding effects, both because making one person richer enables them to be more economically productive and because that person spends some of their money on assets, which themselves generate a return. Moreover, it's been argued that the rate of this social return is much greater than the rate that one can get by investing in the stock market. The existing studies suggest annual investment returns in the order of 7%-48% for iron roofs,

¹⁶ <http://blog.givewell.org/2011/12/20/give-now-or-give-later/>

30%-39% for unconditional cash transfers, 18%-21% for conditional cash transfers¹⁷ and 15-30% for donations targeting contagious diseases.¹⁸

If this were the correct analysis, then it would provide a strong argument for giving now. However, this analysis is not correct. If you were investing to give later, then you would reinvest any financial returns you got in whatever was the highest-return investment. In contrast, you do not have control over the additional money generated by the social return; that money is in the hands of the beneficiaries, who are very unlikely to reinvest their gains in whatever has the highest social return. Instead, the positive effects of your donation spread out more and more over time, until eventually those effects will grow at the rate of world economic growth. The world growth rate, at about 2.25% per year, is lower than the rate of return that you can get from your investment.

Because of this, it doesn't make sense to think of social return as something that compounds over time. Though you should think of the indirect benefits of your donation as increasing the total value of your donation, you should think of it as a fixed benefit. The social rate of return doesn't provide an argument in favor of giving now, because it doesn't matter when the donation happens; you'll get similar indirect benefits from the donation either way.

However, the social rate of return might provide a consideration in favor of giving earlier if it's true that the *rate* of the social rate of return is going down over time. This may be true; if so then, in effect, it adds another wrinkle to the idea that the best giving opportunities are disappearing over time. If, for example, the best investment opportunities for the very poorest are drying up over time, and the multiplier effect is a large component of the total good that you do through your donation, that may mean that the most cost-effective opportunities are disappearing even faster than you might otherwise think.

It's also worth noting that, for some cause areas, there is not social compounding of the benefits of one's donation at all. If one donates to mitigate species loss, or protect the natural environment, or to rescue animals from factory farms, then there may be no social compounding effect.

So far, I've suggested that the idea of social rate of return does not provide an argument either way for giving now or later. This crucially relied on the assumption that the returns from the donor's giving are not captured by the donor, so they can't be reinvested in the best giving opportunities.

Interestingly, however, there is *one* sort of social return where you can 'capture' the donation. This can happen if an organization is dedicated to fundraising (a 'meta-

¹⁷ <http://www.givewell.org/international/technical/programs/cash-transfers#Whatreturnoninvestmentdocashtransferrecipientsearn>

¹⁸ Molyneux, David H. "'Neglected' diseases but unrecognised successes—challenges and opportunities for infectious disease control." *The Lancet* 364, no. 9431 (2004): 380-383.

charity’) and raises at least \$1 with every \$1 it spends on fundraising. If so, then the organization functions like an investment vehicle, but where the money can often get a far higher return than it can in a typical financial investment. For example, Giving What We Can (disclaimer: which I cofounded and run), is an organisation that encourages people to pledge to give at least 10% of their income to the most cost-effective charities, and has over its lifetime raised millions of dollars for those charities. Its very lowest-bound estimate is that every \$1 donated to Giving What We Can results in \$6 donated to top-recommended charities (its best-guess estimate is that every \$1 donated to it results in \$100 of value to top-recommended charities).¹⁹ If the money that it raised were simply directed to global poverty charities, then the fact that it acts as a multiplier on your donation would not give an argument for giving now rather than later. However, a significant proportion of the money it raises is subsequently donated back to Giving What We Can, which can then be used for further fundraising. As long as you believe that the donors that Giving What We Can raises money from have similar values and worldviews to you, or even if not will donate to similar charities that you would choose to donate to, then you are in effect gaining more money to donate later. You might not possess that money yourself, but it will be used in a similar way to how you wanted to use it.

Of course, this argument relies on there being such opportunities available. If donating to a meta-charity really can get such a great return, then many other people will also want to donate to that charity, and the available room for more funding may get used up.

III. A qualitative framework

In the previous two sections, we canvassed a number of considerations that impacted the decision to give now or save and give later. We found considerations relating to tax efficiency, future weakness of will, neglectedness, uncertainty and self-interest were fairly minor, but considerations relating to special obligations, changes in opportunities, knowledge and values, and investment returns were more major.

In this section I propose a qualitative framework that encodes the major considerations, as follows:

Consideration	Score	Score Explanation
Special relationships		1 = special relationships to people now 4 = no special relationships 7 = special relationships to people in the future
Changing opportunities		1 = better opportunities now 4 = no change in quality of

¹⁹ <https://www.givingwhatwecan.org/impact/>

		opportunities 7 = better opportunities in the future
Knowledge		1 = better knowledge now 4 = no change in knowledge 7 = better knowledge in the future
Values		1 = better values now 4 = no change in quality of values 7 = better values in the future
Movement-growth		1 = donation creates many more like-minded donors 3 = no effect 7 = donation decreases the number of like-minded donors
Amount of money to donate		1 = total real value of your money to donate will significantly decrease over time 4 = total real value of your money to donate will stay the same over time 7 = total real value of your money to donate will significantly increase over time

In this framework, I propose that under “amount of money to donate” we bundle considerations relating to taxes, weakness of will and uncertainty as well as financial investment. Each of these considerations simply affects how large the total amount of money to be donated is; whereas the other considerations, in different ways, affect how valuable the expected donation opportunity is. That’s why the idea that the total real value of money assigned to donations might decrease over time is not an idle one: if weakness of will is a real concern (and a DAF is not an option), or if it’s more tax efficient to give now (and again a DAF is not an option), or if one is considering the question at an unfortunate point in time (such as during hyperinflation), then one would, in effect, have more money to donate now rather than later.

On this framework, a total score of 24 would indicate indifference to giving now versus giving later. A score of less than 24 would indicate that in general the arguments favor

giving now. A score of greater than 24 would indicate that in general the arguments favor giving later.

However, because this framework is purely qualitative you shouldn't put much weight on the total scores: the framework should be thought of as a way of thinking through the relevant considerations, rather than giving you a determinate answer. If, for example, we'd carved "amount of money to donate" as several different considerations — "investment opportunities", "weakness of the will", "tax efficiency" and so on — then, because each consideration is currently equally weighted, we would start weighing the "amount of money to donate" consideration several times as highly as we do in my framework.

It is my hope that future work can build on this chapter and develop a quantitative model, which could, at least in principle, give a determinate answer to the question of whether one ought to give now or later, for a given set of inputs. For the time being, I'd encourage each user of the framework to assign their own weights to each of the relevant considerations. If, for example, you think that values change is a particularly important factor, then you should weight that consideration more heavily than the other considerations.

If we were to make this framework quantitative, what might we find about the weights of the different considerations? It's plausible to me that the final term — amount of money to donate — is actually the least important part of the equation, even though it's the one that tends to get the most attention in discussion of giving now versus later (perhaps because the most quantifiable). Even using a 7% real rate of return, over thirty years, investing one's money increases the value of one's donation by only 7.5x. At 5% real return over thirty years, the increase in financial value of one's donations is only 4.3x. However, changes in the expected cost-effectiveness of one's donations, whether through changing opportunities, knowledge, or values, can be much larger. Let me illustrate.

Opportunities and knowledge

It's plausible that the differences in cost-effectiveness between social programs can vary by a factor of tens or hundreds; paying for bednet distribution via a donation to Against Malaria Foundation is estimated to have a cost per life saved of \$3500,²⁰ whereas the marginal expenditure to save a life in the US by the government is about \$7 million.²¹ If we value all lives equally, then, for someone who is initially focused on lifesaving programs in the US, the value of the information that they could instead save lives by funding bednets is vast. And if there is such a large discrepancy between domestic lifesaving programs and international lifesaving programs, then perhaps the discrepancy between global health and some other cause area might also be very great.

²⁰ <http://www.givewell.org/charities/against-malaria-foundation>

²¹ MacAskill, William. *Doing Good Better: Effective Altruism and a Radical New Way to Make a Difference*. Guardian Faber Publishing, 2015.

It's difficult to know how likely it is that we might make such discoveries. But our understanding of the most cost-effective programs is still in its infancy, the result of only a couple of decades of work. It's not implausible to me that we could very significantly improve our knowledge over time.

Special relationships and values

The additional weight you give to the interests of those you have special obligations to is an important factor in the equation. Non-consequentialist views on which special relationships have moral value can easily weight the interests of family, friends, or co-nationals one hundred times as much or more as the interests of distant strangers. The ability to spend money to help those to whom one has a special obligation would therefore matter far more than the additional total amount of money one could donate if one saved in order to donate.

The potential for one's values to change over time is similarly important. For example, suppose that you currently value the interests of co-nationals one hundred times as heavily as the interests of people in other countries; for that reason, you think that the best domestic giving opportunities are about as cost-effective as the best giving opportunities to help the very poorest people in the world. However, suppose that you were to change your mind, and regard the interests of people living in other countries to be just as important as the interests of people living in one's own country. In which case, the best giving opportunities would now provide one hundred times as much value, with a given amount of resources, as you thought before.

Or suppose that you currently endorse a person-affecting view of population ethics, on which there is no value to creating new flourishing lives. For that reason, you think that global poverty is a more important area to focus on than mitigating the risk of human extinction. However, if you were to change your view and adopt the total view of population ethics, then you would come to believe that efforts to mitigate extinction risk were more important; moreover, because human extinction would involve the loss of hundreds of trillions of lives in the future, efforts to mitigate human extinction now, given that the total view is correct, would provide hundreds or thousands of times more value with a given amount of resources than efforts to improve the lives of people in poverty.

Again, therefore, the potential discrepancy in the amount of money one can donate over time is small compared to the possibility that you will rationally change your moral views.

Conclusion

In this chapter I canvassed all the significant considerations that I know of that are relevant to the question of whether to give now or later. The most important

considerations, in my view, are: whether one has the opportunity to invest in movement-growth, and how good that opportunity is; the extent to which the amount of resources flowing into one's favored cause is increasing; to what extent one should expect one's ability to pick the most effective programs will improve over time; and to what extent one should expect one's values to improve over time.

Further work in this area would: (i) develop the qualitative framework that I have proposed into a quantitative framework; (ii) apply this framework to a variety of different cause-areas; and (iii) apply this framework to the *cause-neutral* donor — someone who is open to donating to any problem, and simply wishes to maximize the amount of good they do.