

Being right on the money

Hilary Greaves, Faculty of Philosophy, University of Oxford

'Effective altruism' is defined at www.centreforeffectivealtruism.org as "a growing social movement founded on the desire to make the world as good a place as it can be, the use of evidence and reason to find out how to do so, and the audacity to actually try." One means of trying to do good that is of central concern to effective altruists – by no means the only one, but the one that will be the focus of this essay – is donating money to carefully selected charities. Within this project, the key practical questions faced by any would-be effective altruist are (1) *how much* money to donate, and (2) *which* organisations to donate it to. Consistent with its guiding desire to make the world better, the effective altruism movement generally encourages (1) giving *more* money, and (2) giving whatever money one does give to *more cost-effective* charities. One charity is more cost-effective than another if the first does more good than the second per dollar donated, where 'more good' is understood in a way that, to the greatest extent possible and with due recognition of the fact that not all comparisons can be anything like precise, draws judicious comparisons even between quite disparate cause areas (for example, health vs education, or either of those vs climate change mitigation).

Practical decisions, especially ones like this, are informed in part by one's moral beliefs. Here, the relevant moral issues concern where the line lies, both in the case of how much to give and in the case of where to give, between what one is *morally required* to do on the one hand, and what would constitute going beyond the call of duty on the other. If one fails to meet the requirements of morality, one acts in a way that is morally wrong. By contrast, to go beyond the call of duty is to act in a way that is *better* than barely meeting the moral requirements, but one does not act *wrongly* if one declines to 'go above and beyond'. And while most of us – in particular, most of us who find the idea of effective altruism at all appealing in the first place – feel some motivational pull towards doing things that make the world better simply because they make the world better, most of us also feel a stronger pull towards avoiding moral wrongdoing. Where, then, does the line of moral requirement lie?

Maximising utilitarians, those staples of undergraduate ethics textbooks, give extremely demanding answers to this question. According to maximisation, one is morally required to act in whatever way will bring about the best available state of affairs; according to utilitarianism, the best state of affairs is the one involving the greatest total amount of well-being, summed over all people, weighting the happiness of each person equally. Putting these two commitments together, it seems to follow, in the first instance, that one is morally required to keep giving away money until one is as poor as the world's poorest: for until and unless that point has been reached, giving the diminishing contributions that additional money generally makes to welfare as one gets richer, one could always generate a higher total amount of welfare by giving an extra penny to someone poorer than oneself. And secondly, it seems to follow that one is morally required to give all one's donations to the world's most cost-effective charities, and to those alone: for any time one gives any money to a less cost-effective charity, by definition, one could have achieved more good in the world by giving that same sum of money to a more cost-effective charity.

These answers, however, are deeply unsettling. The familiar problem with a moral demand to give away almost all of one's money is that that seems *too* demanding. Nobody actually does this, and

nor does it take a pessimist to concede that there is no realistic prospect of anyone doing this any time soon. If this were what we really thought morality required, most of us would quickly lose our strong pull towards avoiding wrongdoing, and simply learn to live with our shared status as routine violators of the highfaluting but practically irrelevant moral code.

The problem with a moral requirement to give exclusively to cost-effective charities is less glaring, but on reflection, it too is unsettling. On a fairly regular basis, for instance, some neighbourhood schoolchild knocks on my front door, reporting that he or she is doing (say) a sponsored swim for (say) Save the Children, and asking for my support to the tune of a few pounds by way of sponsorship. Now, Save the Children, like the vast majority of more-or-less randomly selected charities, does not top the cost-effectiveness charts, for the simple reason that only a tiny minority of charities can be top of the charts. But what am I supposed to do? To say, or even to think, that I shouldn't sponsor my neighbour's child, on the grounds that my £5 would do even more good if donated to a cost-effectiveness chart-topper, seems churlish.

That is a case in which I am motivated to give to a less cost-effective charity by some personal connection to a *benefactor* (my neighbour's child). Other relevantly similar types of cases arise when I have some personal connection either to the charity's *beneficiaries*, or to the *cause area* in which the charity works. While travelling I might, for example, spend some time in a refugee camp, chatting with the camp's inhabitants and organisers and witnessing at first-hand these people's dire need for food, clothing, medical supplies and educational resources. I could offer to dig into my pocket and help – or I could keep quiet and keep my money, in order to donate it later to the chart-toppers. While there is a clear story about why the second course of action is rational and moral, here too one would feel churlish.

Cases of personal connection to a particular cause area are, if anything, even more ubiquitous. If I have personal experience of sexism, that might well motivate me to give selectively to charities whose work is focussed on opportunities specifically for women and girls. Or if I have recently lost a family member to cancer, that might well motivate me to donate to charities focussing on cancer research or care. But here, too, it is statistically very unlikely that there will be any charity in the specific area in question that happens to top the cost-effectiveness charts. No cancer charity, for example, comes anywhere near to the Against Malaria Foundation in terms of estimated cost-effectiveness on widely agreed metrics.

If we are unsettled by the very demanding answers to our questions about (1) how much money we are morally required to give and (2) what are the moral requirements concerning where we give it to, we might consider more permissive answers. A permissive answer to the "how much?" question might say that one is only *morally required* to give (say) 10% of one's income, while acknowledging that giving more than that would be *better*. And a permissive answer to the "where?" question might say that one is only morally required to direct one's donations to (say) organisations that are at least moderately cost-effective compared to others operating in the same cause area, while acknowledging that giving to more cost-effective charities instead would be better.

While public discussions of these issues in effective altruist circles relatively rarely explicitly use the language of moral requirement, there is arguably a general sense in the effective altruist air that favours a permissive answer to the "how much?" question, but sticks to a demanding answer on the "where?" question. Insofar as there is, this is entirely understandable *from the point of view of*

marketing strategy. In the first instance, as has been repeatedly emphasised, the difference between more vs less cost-effective charities is not merely a matter of a factor of two or three, but often a matter of two or three orders of magnitude: so, in general, especially if one is willing to shift between cause areas in response to evidence on cost-effectiveness, *where* one donates will likely have a far greater impact than *how much* one donates. And secondly, shifting one's donation to a more cost-effective cause area is generally far less costly to the donor than upping the amount of one's donations, so that the arena of cost-effectiveness is one in which there is less conflict between impartial good achieved vs narrow donor self-interest. 'Adverts' for effective altruist behaviour are therefore on considerably more persuasive ground, for most people, when they target questions of where to give rather than questions of how much to give.

Marketing strategy, however, is one thing, and moral truth quite another. And when we turn to the reasons that one might give, at the level of principled moral theory, for a not-too-demanding answer to the "how much?" question, it is actually remarkably difficult to find any such reasons that do not equally rationalise a more permissive answer to the "where?" question.

Clearly, any moral theory that gives a permissive answer to *either* of our key questions must deviate from the austere maximising utilitarianism we sketched above. The question is what the deviation in question is. This is not the place to attempt a comprehensive survey of possible moral theories. But at least one very natural theory, itself not entirely opposed to the basic spirit of effective altruism, recognises "agent-centred prerogatives" alongside a basically utilitarian recognition of the equal value, from the "point of view of the universe", of each person's happiness. According to this theory, considerations of impartially totalled-up welfare constitute *one* important factor in determining what, in any given situation, one is morally required (or permitted) to do; but this is not the only factor. In addition to this impartial point of view of the universe, there is also the agent's own personal point of view, from which things get weighed up differently. In the personal point of view, the agent's own interests and projects, and those of her friends and family, are given a special extra weighting, out of proportion to their impartial significance, simply because they are the agent's *own* (interests, projects, friends or family). Concerning the notion of moral requirement, the idea is then that one is not in general morally required to maximise the impartial good: one is permitted instead to take actions that lead to significantly less than the maximum attainable amount of impartial good, when such 'failures to optimise' are rationalisable in terms of an increased amount of "agent-centred" good, non-impartial goodness as it appears from the agent's own personal point of view.

This alternative theory – just by way of example – arguably has a ready answer to the question of how come we are not morally required to give away almost all of our money. We are not morally required to do that, this theory answers, because although doing that (or something close) would perhaps lead to the greatest amount of good *considered from the impartial point of view*, it would lead to significantly less agent-centred good. Both types of good are relevant to questions of moral requirement; considerations of impartial good are certainly important, and accordingly we are morally required (this theory says) to give quite a bit, perhaps, but nothing like the amount the maximising utilitarian would have us believe.

If that, or anything like that, is our rationale for a permissive answer to the "how much?" question, however, precisely the same set of theoretical considerations naturally generates a rationale for a permissive answer to the "where?" question. For, in the examples of less than maximally cost-

effective giving mentioned above, the agent's own interests and projects significantly favour a less cost-effective over a more cost-effective charity. It is psychologically at least somewhat costly for me to say no to the child on my doorstep, even if I am doing so in order to donate the money in question to an impartially better charity. And by definition it is costly, from the non-impartial point of view of my own concerns and projects, to hold back from giving anything to causes that are closer to my heart – the refugees that I have formed relationships with, the causes of feminism or cancer that I have particular personal reasons to care about – simply on the grounds that, considered from an *impartial* point of view, the same money could do more good elsewhere. More generally, it seems difficult to come up with any halfway sensible theory that rationalises a permissive answer to the “how much?” question, without equally rationalising at least a somewhat permissive answer to the “where?” question. Perhaps we are morally required to give *most* of our charitable donations to the cost-effectiveness chart-toppers, but, if we're morally permitted to keep any appreciable portion of our income ‘for ourselves’, it's hard to see how morality could then forbid us from instead donating some of it to less than maximally cost-effective charities, when we have a good personal reason to want to.

When we have a good personal reason to want to. The qualification is important. Nothing I have said here suggests any rationale for, or defence of, charitable giving that is essentially random. This is the bizarre, but all-too-common, giving behaviour in which one gives to charity more or less on a whim, without any thought for, much less research into, the cost-effectiveness of the charity concerned. Granted that I have a particular personal connection to the cause areas of (say) cancer care and research, still what I want is for my cancer-related donations to do as much good as possible for cancer patients, per pound I donate. I will not give to Cancer Charity 1 if the evidence suggests that the same donation could do ten or a hundred times as much good for a relevantly identical group of cancer patients if donated to the more effective Cancer Charity 2, whether the reasons for the difference in cost-effectiveness are matters of the tractability of cases focussed on by the two charities, over- or (more likely) under-spending on CEO pay and administrative overheads, or anything else. Insofar as my motivation for charitable giving is the acquisition of warm fuzzies for myself rather than pure altruistic concern, I will not even get many warm fuzzies if I give to Cancer Charity 1 having failed to check whether giving £100 to Cancer Charity 2 actually does more good for cancer patients than giving £1000 to Cancer Charity 1 – I will just feel like a fool. And if the amount of my donation is beyond the merely trifling, certainly I will regard it as worth the time investment to do a bit of cost-effectiveness research, given that, in this age of charity evaluators, all that this actually requires is checking the website of some suitable independent evaluator, such as GiveWell or Giving What We Can. I would be crazy not to do this, just as I would be crazy not get at least two or three independent quotes before spending several thousands of pounds on building works for my house. This part of effective altruist logic seems unassailable, and the behaviour that ignores it does seem, if not actually immoral, at any rate irrational.