Risky Giving[1]

Theron Pummer
University of St Andrews

Peter Singer famously argues that it is wrong of us not to give to effective charities helping people in extreme poverty. In "Famine, Affluence, in Morality" (1972), he offers an elegant argument for this conclusion.

The *ethical premise* of his argument is that, if we can prevent something very bad from happening, without thereby sacrificing anything nearly as morally significant, then it would be wrong not to do so. In case this premise isn't obvious enough on its own, Singer offers an example: suppose you're walking past a shallow pond and you see a child drowning in it; you realize that the only way to save the child's life is by rushing in quickly, ruining your clothes and shoes. Virtually everyone agrees that it would be wrong of you not to save the child. What seems relevant in this case is that you can prevent something very bad from happening (the child drowning) without sacrificing anything nearly as morally significant (the loss of your clothes and shoes isn't nearly as morally significant as the loss of this child's life). Arguably, given that we think it's wrong to let the child drown, we should also accept the ethical premise of Singer's argument.

The *empirical premise* of his argument is that in fact we *can* prevent very bad things from happening, without thereby sacrificing anything nearly as morally significant, by giving to effective charities helping people in extreme poverty. So, it is wrong of us not to give to such charities.

One of the most impressive features of Singer's argument is that it builds from premises that look plausible, independently of the broader ethical theory, if any, one accepts. Singer himself accepts *consequentialism*, the ethical theory according to which an act is wrong if it fails to bring about the overall best outcome available to the agent. But many philosophers and non-philosophers alike instead accept *non-consequentialism*, and recognize the existence of *moral options* and *moral constraints*. Moral options entail that it is sometimes *not* wrong for an agent to do less than what would result in the overall best outcome. For example, even if it would make the world a better place if you gave up your kidney to a stranger in need, intuitively it would not be wrong of you not to. Moral constraints entail that it is sometimes *wrong* for an agent to do what would result in the overall best outcome. For example, even if it would make the world a better place if you *stole* someone's kidney and gave it to someone else in need, intuitively this would be wrong. It isn't that non-consequentialists deny the moral importance of making the world a better place, or helping others; what they deny is that this is the *only* thing that is fundamentally morally important.

Given the entrenched disagreement between consequentialists and non-consequentialists, an argument for giving to charity that takes consequentialism for granted would have limited appeal. But the ethical premise of Singer's argument is compatible with consequentialism as well as non-consequentialism. The premise leaves unspecified what counts as "nearly as morally significant" as preventing a very bad thing from happening. Bringing about something *else* that's very bad, or failing to bring about something that's very good, could presumably be roughly as morally significant as preventing a very bad thing. But

arguably doing something that's substantially costly to you or to those near and dear to you, or doing something against which there is plausibly a moral constraint (such as harming, killing, stealing, lying, promise-breaking, and so on) can also be roughly as morally significant as preventing a very bad thing. The ethical premise would not imply that it would be wrong to refuse to give until you're poor yourself, as – from the perspective that you have the moral option not to do what results in the best outcome – that is arguably as morally significant as preventing much greater bad things. The ethical premise would also not imply that you have a moral obligation to steal from the rich and give the money to charity, as – from the perspective that you are under a moral constraint not to steal even if this results in the best outcome – that is arguably as morally significant as preventing much greater bad things. Singer intentionally leaves it quite unspecified what counts as "nearly as morally significant" in order to give his argument *wide appeal*. On any serious, minimally reasonable way of filling in "nearly as morally significant," the ethical premise will appeal to consequentialists and non-consequentialists alike as well as work together with the empirical premise to deliver Singer's famous conclusion.

One might fail to be convinced of Singer's argument on empirical grounds. In particular, one might doubt that in fact we can prevent very bad things from happening by giving to effective charities helping people in extreme poverty. One might doubt whether there are any such *effective* charities.

The causal processes by which your money does things for those in extreme poverty are, I am afraid, far more complex than most charities let on. Worse still, taking the bad effects of charities together with the good ones, it is for very many charities extremely unclear whether your donations are likely to prevent more bad than they produce. While experts on government aid for international development have been around for decades, until fairly recently one would have been hard-pressed to find anyone who might reasonably qualify as an expert on *individual giving*: someone who has carefully and systematically empirically studied the net impact per dollar that everyday individual donors can have by giving to various charities. Of course, we've long been exposed to the claims of impact embedded in advertisements from the charities themselves, but, to put it gently, these aren't always the most unbiased sources of information. Many of us have come across websites (like *Charity Navigator*) that impartially rank charities according to what proportion of their donations they spend on administrative costs rather than their actual programs. But in the absence of information about the likely *effectiveness* of these programs, such websites are pretty pointless as a guide to individual giving.

Fortunately, for almost ten years now, two independent, transparent, rigorous, and reputable *charity evaluators* have come on the scene. These are *GiveWell* and *Giving What We Can*. After reviewing thousands of different charities, these organizations together recommend only about half a dozen. It is worth noting that both organizations have ranked *Against Malaria Foundation* (which distributes insecticide-treated bednets) and *Schistosomiasis Control Initiative* (which distributes deworming tablets) in their highest categories. In light of the convergence among the experts on individual giving, I believe that the everyday individual donor can be reasonably confident that donations to these and similar charities are very likely to prevent much more bad than they produce. Indeed, impressed with the amount of good individuals can do with relatively small contributions, and with the enormous differences in cost-effectiveness between different charities, many now subscribe to the philanthropic philosophy and social movement known as *effective altruism* (see other contributions to this symposium).

There is excellent evidence of the existence of effective charities helping people in extreme poverty, which supports Singer's empirical premise. However, one might doubt the empirical premise in another way: one might doubt that in fact we can prevent very bad

things from happening, *without thereby sacrificing anything nearly as morally significant*, by giving to effective charities helping people in extreme poverty. In particular, one might worry that in giving to charity, unlike rescuing the child from the shallow pond, you significantly *risk harming* individual people. In his essay "Poverty is No Pond" (2010), Leif Wenar writes (p. 117):

> From the perspective of an individual living in poverty, a particular aid effort may harm more than it helps. The aid may strengthen the autocrats, corrupt bureaucrats, warlords, soldiers or criminals that have power over that person. A humanitarian effort may draw the person away from their self-sustaining livelihood; a development project may draw them into acquiring skills for which there is no employment. Aid in aggregate may increase inflation, reduce employment, or weaken the provision of public services. Aid flowing into this person's country may delay needed political reforms, and make both the government and the citizens more responsive to foreigners than to each other. Aid may inflame economic inequalities or ethnic antagonisms in ways that are bad for this person, or damage his or her self-esteem. And so on. Even assuming an aid project or official aid to a country helps overall, it may leave at least some people worse off than before.

Wenar spells out each of these risks in detail in his essay, which I highly recommend reading. In the case of the drowning child, you can *allow* harm to befall an individual, leaving the child as badly off as she'd be if you did nothing. In the case of giving to charity, on the other hand, not only can you allow harm to occur, but you can also *impose* harms on individuals, making these particular individuals worse off than they'd be if you did nothing. The latter is of special moral concern to non-consequentialists. While it's wrong to fail to save a drowning child, intuitively it is morally worse to grab a child off her tricycle and *drown* her (even assuming the outcome is the same either way).

Consequentialists believe it is morally permissible to impose risks of harm on individuals, provided that it's part of what's likely to be best overall. Many non-consequentialists do not. It is not generally morally acceptable, on their view, to impose harms on some people, even if *other* people benefit enough such that things are better overall. Benefits to some do not straightforwardly morally cancel out the harms imposed on separate persons (e.g., it is not the case that drowning one child and rescuing another is morally equivalent to doing neither). One might object that Singer's argument fails to be convincing to non-consequentialists, despite his efforts. In particular, one might claim that, as long as the notion of "nearly as morally significant" leaves room for moral constraints against risking harm to individuals, the empirical premise will fail to obtain (however plausible the ethical premise is).

I believe this is a mistake, and that such concerns about risks of harming individuals should not deter non-consequentialists from embracing Singer's conclusion. The short version is that in some realistic cases of giving to charity, it is not wrong to impose these risks; typically they are not aimed at in order to achieve further ends but instead are generated as side effects of interventions that are likely to do a lot of good overall, and often they are distributed across people in a way that is morally acceptable.

In their important book *Responding to Global Poverty* (Cambridge University Press, 2016, p. 132) Barry and Øverland claim that it can be morally acceptable to impose a risk of harm on someone if doing so is part of providing *that same person* with a chance of benefit, such that, taking into account the likelihood of the harm as well as the benefit, things are on balance likely to be better for this person. Call such balancing *intrapersonal risk balancing*. On the other hand, consider imposing a risk of harm on someone as part of providing

someone else with a chance of benefit, even if things are on balance likely to be better overall. Call such balancing *interpersonal risk balancing*. According to many non-consequentialists, this latter type of risk balancing is other things equal more morally objectionable than intrapersonal risk balancing. To get a more concrete handle on these two types of risk balancing, let's have a look at a pair of imaginary cases.

The first is called *Risks to the Ill*: One hundred people are dying of a terrible illness. If left alone, they'd each have only a 20% chance of survival. However, they can take drug X, which would increase their chances of survival to 80% each. The only way to get drug X to them is by putting it in the water supply, and, we can suppose, there are no other people nearby to drink this water. However, drug X has a negative side effect. It risks liver damage, imposing a 10% chance of death on each of the one hundred ill people who take it (absent this risk taking drug X would increase their chances of survival to 90% each).

Now let's consider a second case, called *Risks to the Healthy*: One hundred people are dying of a terrible illness. If left alone, they'd each have only a 20% chance of survival. However, they can take drug Y, which would increase their chances of survival to 90% each. The only way to get drug Y to them is by putting it in the water supply, which, we can suppose, would mean one hundred healthy people would also receive drug Y. And, while drug Y has no negative side effects when ingested by people with the illness, it risks liver damage when ingested by healthy people, imposing a 10% chance of death on each of them.

Risks to the Ill is an example of intrapersonal risk balancing, whereas Risks to the Healthy is an example of interpersonal risk balancing. Other things being equal, intrapersonal risk balancing seems morally acceptable given that each person is likely to be better off on balance. Interpersonal risk balancing seems more morally problematic, even if it is likely to be best overall. It's not that non-consequentialists must be absolutists about this, claiming that it is always wrong to engage in interpersonal risk balancing, no matter how small the risks to others are and no matter how good things will likely be overall. That would be implausibly extreme, and it would imply that virtually all of our actions wrong. No matter what you do, there is always *some* risk that through some butterfly effect an individual elsewhere will be harmed. Instead, what the non-consequentialist should say is that interpersonal risk balancing *can* be morally justified, but that to the extent that the risks to others are large, and to the extent that the harms they might incur are large, the imposition of such risks would have to make things likely to go *proportionally much* better overall, to be morally justified. And perhaps when the risks and associated harms are large enough, some non-consequentialists would be attracted to a kind of absolutist view, claiming that such interpersonal balancing is wrong, no matter how good overall the effects are likely to be.

The imaginary cases described are unrealistic in the following respect: Risks to the Healthy is a *purely* interpersonal risk balancing case, and Risks to the Ill is a *purely* intrapersonal risk balancing case. In the real world, often cases will be *mixed*, being more or less like one of the pure cases.

There is a further complication. In intrapersonal risk balancing cases like Risks to the Ill, each of those who would receive the risks of harm alongside the chances of benefit is on balance likely to be better off. But suppose that, on reflection, these ill people prefer not to have the intervention simultaneously bringing risk of harm and chance of benefit. Then it may seem that nonetheless imposing the intervention on them would be disrespectful and morally problematic. But what if only one out of the one hundred in Risks to the Ill prefers not to have the intervention, while the other ninety-nine very much do prefer it? Again, a non-consequentialist need not be an absolutist about this, claiming that it's wrong to engage in intrapersonal risk balancing whenever *anyone* involved prefers not to have the intervention, no matter how many others prefer it and no matter how good things will likely be overall. That would be implausibly extreme, and it would imply that virtually all policy

decisions are wrong. For any policy decision, it is likely that at least one person affected by it will prefer not to be so affected. Instead, what the non-consequentialist should say is that intrapersonal risk balancing *can* be morally justified even when it is not preferred by some of those affected, if enough others affected do prefer it and if it makes things likely to go sufficiently better overall.

What all this suggests is that, when our charitable giving would impose risks of harm on particular individuals, there are a number of things we need to know to determine whether it would be wrong to proceed. First, we need to know how good overall each of our acts is likely to be. According to consequentialists, this is all we need to know. But according to non-consequentialists who believe there are moral constraints on imposing risks of harm on individuals, we need to know more. To what extent would we be engaged in intrapersonal rather than interpersonal risk balancing? (Again, realistic cases will often involve a mixture of the two types of risk balancing.) How serious are these risks to individuals, in terms of degree of likelihood and degree of harm? To what extent do those affected by our act prefer the package of risks of harms and chances of benefits it delivers to them?

I believe that several effective charities do quite well in terms of these non-consequentialist criteria, in addition to the consequentialist one. You can assess this for yourself. Visit the GiveWell website (givewell.org) and the Giving What We Can website (givingwhatwecan.org). Look up their top charities, particularly Against Malaria Foundation (AMF) and Schistosomiasis Control Initiative (SCI). Read their detailed reports and follow links to other documents where appropriate, etc. In addition to studying the positive effects of these charities, be sure to have a look at GiveWell's analysis on "possible negative or offsetting impact" (e.g., they discuss the risk that distributing bednets in certain areas will lead to overfishing, if they are misused). Consider this empirical data on AMF and SCI from GiveWell and Giving What We Can against the backdrop of the general risks of giving to overseas charities, discussed by Wenar. Seek out and listen to the views of experienced people who have worked in development overseas.

Giving to AMF or SCI is certainly not *risk-free*, but in my view it involves relatively little risk of doing serious harm. Moreover, it is important to keep track of *which particular people* these risks of harm accrue to. To the extent that the risks of harm are imposed strictly on those immediately affected by AMF, for example, there may be no non-consequentialist objection here at all. This is because the likely benefits to those immediately affected by AMF's operations are large enough that it'd be plausible to classify this as a purely intrapersonal risk balancing case. But going back over Wenar's list of risks of harm, one might observe that these risks pertain to people well beyond those immediately affected by the given charity's operations. Indeed, nearly everyone within the country or broader geographic region of the charity's operations is potentially exposed to these sorts of risks. And it is harder to morally justify imposing such risks on people who do not at the same time enjoy chances of benefiting.

One of the particularly brilliant things about charities like AMF and SCI, however, is that they bring with them significant *chances of benefits* to these wider groups, too. GiveWell and Giving What We Can cite research suggesting that fighting malaria and neglected tropical diseases are likely to have positive economic effects, which extend to everyone within the country or broader geographic region of these interventions. This fact makes giving to AMF or SCI *more like* a case of intrapersonal risk balancing, and so more morally acceptable. To the extent that these chances of benefits to each of those in the wider region outweigh the risks of harms to each of those in this wider region, we'd again have a case of purely intrapersonal risk balancing, which is morally acceptable. A consequentialist might view such chances of wider benefits as merely another factor increasing the likelihood of making things better overall. But to a non-consequentialist concerned about the morality of

imposing risk of harm on particular individuals, they can serve to partially or fully remove what otherwise would have been a moral obstacle to giving. All things considered, my view is that AMF and SCI do well in terms of the non-consequentialist criteria outlined above, and that we can prevent very bad things from happening, without thereby sacrificing anything nearly as morally significant, by giving to them.

Not everyone will, upon studying the empirical data, come to the same conclusion I have. But even if we should not give to any overseas charities, on grounds of risk of harm, does this mean that Singer's conclusion is false? Not exactly, as we could save up and grow our money rather than give it away *now*; we could instead give at a later time when presumably the risks of giving to charities are better understood. Of course, this carries with it the risk that we will never actually give! Another option is to give to promote charity research (for example by giving to GiveWell's Open Philanthropy Project), so that our donations can be used to accelerate progress on understanding the likely benefits and harms of giving to charities helping those in extreme poverty. This form of giving carries virtually no risk of harming others, and it is likely to do considerable good in the long run. This is another way of preventing very bad things from happening, without thereby sacrificing anything nearly as morally significant.