

Aid Scepticism and Effective Altruism

WILLIAM MACASKILL

University of Oxford

ABSTRACT

In the article, ‘Being Good in a World of Need: Some Empirical Worries and an Uncomfortable Philosophical Possibility,’ Larry Temkin presents some concerns about the possible impact of international aid on the poorest people in the world, suggesting that the nature of the duties of beneficence of the global rich to the global poor are much more murky than some people have made out.

In this article, I’ll respond to Temkin from the perspective of effective altruism—one of the targets he attacks. I’ll argue that Temkin’s critique has little empirical justification, given the conclusions he wants to reach, and is therefore impotent.

SECTION I

Let us begin by discussing the empirical evidence on aid and economic growth. The majority of Temkin’s article discusses the possible negative impacts that foreign aid could have. Towards the end of this article, however, Temkin makes the following striking comment: “Here I sit, comfortably speculating about various *possible* negative effects that aid groups may produce.... I haven’t offered empirical evidence to support the concerns that I have raised.” [emphasis Temkin’s].

This is a surprising approach, to say the least. One might think that, when discussing the ethics of attempting to benefit the global poor—and, especially, when focusing specifically on whether such attempts are effective or not—the vast empirical literature in development economics would be relevant. But Temkin doesn’t discuss this literature, instead preferring to rely almost wholly for his empirical claims on a single chapter by a single author of a single book that was intended for a general audi-

ence.¹ There is no doubt that Angus Deaton is an outstanding economist. But is he so good as to warrant ignoring the thousands of other economists who have dedicated their careers to understanding the effects of aid on those in poor countries?

Temkin makes the following methodological remark in defence of his approach: “As for gathering and assessing empirical data, I must leave that to the social scientists. My job, as a philosopher, is to help identify both empirical and normative issues relevant to our obligations to the needy, which I have done.” Now, it is true that Temkin’s preferred grammatical construction is modal: the words ‘possible’, ‘may’, ‘might’, ‘can’ and their variants collectively appear over 150 times in his article. And, of course, as philosophers we often should reflect upon the merely hypothetical, because of the broader lessons that we can learn from hypothetical cases.

But I find it hard to believe that this is really the spirit in which Temkin’s article is written. His focus is firmly on the very practical question of what we, in the actual world, are obliged to do. And his final sentence is non-modal: “Unfortunately, what one should do in light of [the fact that the well-off who ignore the badly-off are open to serious moral criticism] is much more complex, and murky, than most people have realized.”

Perhaps, instead, it’s the mere *possibility* of doing harm that means we should potentially refrain from donating to development charities? This would make Temkin’s reliance on a single source more justifiable. But it would result in a very weak argument, because there are possible harms from almost every action that we undertake.

For example, almost all the same concerns that Temkin raises for donations would also apply to consumption of goods produced in the developing world by international companies. Perhaps foreign private investment will result in a brain drain away from other industries or from the country itself (if the employees go to work at some other branch of a multinational company). Perhaps those foreign employers will undermine the local governments in order to get their way. Perhaps, even, some of the money that you’d spend on goods produced in the developing world would be used by companies to bribe government officials in order to get a tax break. Is Temkin therefore recommending that we only buy local, too?

In general, laying out what-ifs doesn’t get us very far in the ethics of development. All activities we engage in have some chance of harming others, and some chance of benefiting others. What matters is how great the possible benefits are, how great the possible harms are, and the probabilities of each. Temkin, unfortunately,

1. Deaton (2013, ch.7).

doesn't attempt to address these issues,² and so, from his article alone, we are left none the wiser about whether development efforts are on average good or bad.

But what should we think about the impact of aid on economic development? It seems to me that there are two reasonable positions that one could have, depending on how sceptical one is towards econometric analysis in this area. If one is sympathetic to the value of econometric methods, then the most natural conclusion is that aid has a positive but modest effect on growth. This view is represented by Jonathan Glennie and Andy Sumner, in a Center for Global Development policy paper:

the assertion that aid generally contributes to economic growth, while not proved beyond doubt, is now less contentious in the academic literature than is currently recognised in public policy debate. That is not to say that there is an absolute consensus, nor that there are not important unresolved questions that would need addressing to claim unequivocal proof, but that aid's critics are currently in the academic minority (Glennie and Sumner 2014).

If one places less credence in econometric methods in this context, then one might reasonably be agnostic either way about aid's effects on growth. The sample size of 80 countries is small, and there are major confounds that are difficult to control for. (If receiving more aid is correlated with lower GDP, is that because aid hinders growth, or because richer governments will give more aid to countries that are poorer?) This view is represented by Owen Barder (himself citing David Roodman) in a report to the British House of Lords: "Given the modest volumes of aid, we should not expect an impact on growth which is bright enough to shine through the statistical fog." (Barder 2011).

SECTION II

However, even if one thinks that there is a lack of robust evidence regarding the effect of bilateral aid on economic growth, that should not lead us to a position of agnosticism about whether one can meaningfully improve the lives of the extreme poor. Crucially, the general debate about aid's impact on economic growth has very

2. See Ravallion (2014) for a discussion of this literature in relation to *The Great Escape*. For a selection of the recent literature on aid, see Burnside and Dollar (2000), Rajan and Subramanian (2008), Clemens et al (2012), and Mekasha and Tark (2013).

limited bearing on the sorts of donations that we should be talking about, such as donations to Against Malaria Foundation, one of GiveWell's top-recommended charities. This is for three main reasons.

First, the debate on aid effectiveness (including the work of Angus Deaton and other aid critics like William Easterly and Dambisa Moyo) is focused on bilateral aid, *not* on non-governmental organisations. But the most obvious donation targets for individuals in rich countries are NGOs.

Second, the vast majority of aid scepticism is aid aimed at economic development, rather than at global health. The track records of these two projects are very different. Though the track record of attempts to foster economic growth is arguably unclear, the track record of global health is astonishing. The eradication of smallpox has saved over 60 million lives since 1980 (more lives saved than if we'd achieved world peace in the same time period) and 1/3 of the funding of the eradication effort came from international aid.³ Globally, rates of death from measles, malaria and diarrhoea are down by 70%.⁴ Indeed, even those regarded as aid sceptics are very positive about global health.⁵ Here's a quote from Angus Deaton, from the same book that Temkin relies so heavily on:

Health campaigns, known as "vertical health programs," have been effective in saving millions of lives. Other vertical initiatives include the successful campaign to eliminate smallpox throughout the world; the campaign against river blindness jointly mounted by the World Bank, the Carter Center, WHO, and Merck; and the ongoing— but as yet incomplete— attempt to eliminate polio (Deaton 2013 p.104-5).

Later in the book he states: "There may ... be cases in which aid is doing good, at least on balance. I have already made that case for aid directed toward health." (Ibid p.318) Similarly, here's Bill Easterly, another aid sceptic:

There are well known and striking donor success stories, like the elimination of smallpox, the near-eradication of river blindness and Guinea worm, the spread of

3. For further explanation of this, see MacAskill (2015, ch.3).

4. Global IDEA Scientific Advisory Committee (2004). For further case studies in global health, see <http://millionssaved.cgdev.org/> [Accessed 22/5/19]

5. These points are made at more length at <https://blog.givewell.org/2015/11/06/the-lack-of-controversy-over-well-targeted-aid/> [Accessed 22/5/19]

oral rehydration therapy for treating infant diarrheal diseases, DDT campaigns against malarial mosquitoes (although later halted for environmental reasons), and the success of WHO vaccination programmes against measles and other childhood diseases (Easterly 2009).

In the same post, he summarises his view by commenting that, “even those of us labeled as “aid critics” do not believe aid has been a universal failure. If we give you aid agencies grief on failures, it is because we have seen some successes, and we would like to see more!”

Finally, what is relevant to individuals in rich countries is not the quality of typical aid programs, or even of average aid programs. What really matters is how effective the best aid programs that we can identify are. As an analogy: suppose that a group of people were championing turmeric ingestion as an antidepressant, and pointing to empirical evidence that supported their view (e.g. Ng et al 2017). How much would we learn about the effectiveness of that particular intervention by discussing how often herbal remedies work in general? Not much.

Luckily for us, GiveWell has conducted years of extensive investigation to try to find outstanding global health and development charities. Though it’s impossible to tell whether the charities they’ve found are the ‘best’ charities, they have certainly found charities that do a huge amount to benefit the poor, are well-evidenced, and transparent.

Consider, for example, the Against Malaria Foundation, one of GiveWell’s most highly recommended charities. A summary of the case for its positive impact is as follows:⁶

1. There is highly extensive evidence that using long-lasting insecticide treated bednets (LLIN) decreases the incidence of malaria and therefore of illness (such as anemia and splenomegaly) and death.⁷
2. These positive direct impacts come at very low costs, perhaps as low as just a few thousand dollars per life saved.
3. AMF focuses on countries that have high rates of malaria, and ensures that their distribution partners conduct post-distribution surveys (at 6-month in-

6. See GiveWell (2018a) for a much fuller analysis.

7. Summaries of this evidence can be found in two Cochrane meta-analyses, Lengeler (2004), which reviews 22 randomised controlled trials, and Gamble (2007), which reviews 6 randomised controlled trials) and on GiveWell (2018b).

tervals for a period of 2.5 years) to ensure that LLINs reach their intended destination and are being used properly.

What about possible negative impacts? GiveWell have looked into many of those, too. Are bednets used for alternative purposes, such as fishing—inadvertently depleting fish stocks due to the nets’ small holes—as suggested by a popular *New York Times* op-ed? (Gettleman 2015) GiveWell have addressed this issue (GiveWell 2015), and, because AMF conduct post-distribution surveys, we can be confident that over 80% of nets are hung up 6 months after distribution. The evidence from the *New York Times* op-ed, in contrast, is based on a single study of household nets distributed along Lake Tanganyika, and provides no evidence that this is connected to depletion of fish stocks.

Might funding of AMF just displace efforts to distribute bednets by other actors? This holds to some extent, and should decrease to some extent our estimate of AMF’s cost-effectiveness (and this is taken into account in GiveWell’s analysis). But the bednet gap across sub-Saharan Africa is very large (requiring hundreds of millions of dollars to fill), and there are a number of clear examples of distributions that AMF have funded that would not have happened otherwise.⁸

Does saving lives lead to overpopulation? GiveWell commissioned an extensive report on this question (Roodman 2014). It is probably true that saving lives has a mild effect of increasing population size, but it’s very unclear whether this should be regarded as a good, bad or neutral thing. Indeed, answering this question requires answering notoriously recalcitrant problems in population ethics.

To be clear: We cannot be 100% certain that AMF is doing more good than harm; nor can we be confident about exactly how much good it’s doing (a fact that GiveWell repeatedly emphasised (GiveWell 2017)). Indeed, there are many issues that remain open. LLINs probably do contribute to insecticide resistance, though the malaria community still recommends the use of LLINs (GiveWell 2013).⁹ It’s possible that distributing free bednets undermines local markets for nets because people expect that they will receive them for free (see GiveWell 2012 for discussion). It’s possible that AMF diverts some skilled labour from other areas, so the full costs of bednet distribution are larger than they might otherwise be (this is discussed in GiveWell 2018a).

8. See GiveWell (2018c). For discussion of how GiveWell accounts for funging with other actors in general, see GiveWell (2018d).

9. In addition, two recent randomized controlled trials suggest that next generation nets can successfully mitigate resistance: see Tiono et al (2018) and Protopopoff et al (2018).

We would be naive indeed if we did not appreciate the fact that development efforts operate in an incredibly complex context of multiple interrelated activities.

But we cannot be certain, for any of our actions, that we will do more good than harm. We should be careful to hold development efforts to the same standard we hold other activities: neither blithely assuming that any good intention will result in good actions, nor being so timorous that any possibility of harm results in paralysis. The best we can do is assess the evidence as best we can. In the case of AMF, there is a very strong case that the direct effects are very positive, a well-established track record of similar interventions being very effective, and no solid extant case for there being significant negative long-run effects.

In sum, Temkin provides no evidence at all for the idea that charities like AMF are doing more harm than good. In contrast, GiveWell provides extremely extensive evidence in support of the idea that AMF are doing much more good than harm. Given that Peter Singer, Temkin's primary target, endorses GiveWell, it's surprising that Temkin takes no time at all to engage with the hundreds of thousands of hours of work that GiveWell have invested to make their recommendation, especially when they often deal directly with the worries that Temkin has.

SECTION III

Temkin does anticipate something similar to this response. He asks: "Why can't Deaton simply support Effective Altruism? Instead of claiming that we shouldn't be supporting international aid groups operating directly in the world's poorest regions, why shouldn't Deaton contend, more modestly, that we must be very careful about which aid groups we support, to make sure that they are, indeed, doing more good than harm?" That is: he raises the idea that what is relevant ethically, is not how effective *typical* aid charities are, but instead how effective the best charities we can find are.

In response, Temkin claims that this leads us to an each-we dilemma. He claims that, by supporting those charities that effective altruists promote (such as AMF), we may each individually do the morally best thing, but we still collectively do more harm than good.

But this is clearly a non sequitur. Temkin's each-we dilemma can only get off the ground if he supposes that (i) the total aggregate harm of a single action that aims to help the poor is greater than the total aggregate benefit of that action; and that (ii)

morally, the aggregate of many small harms can never outweigh a small number of large harms (at least, for some magnitude of benefit and harm).

In other work (MacAskill forthcoming) I have suggested that effective altruism, as a matter of definition, should be committed to anonymity (roughly: that the identities of individuals across outcomes do not matter morally), and in this volume Clark and Pummer argue—correctly in my view—that anonymity is incompatible with (ii).

What I have argued in this article is that, for the charities that GiveWell recommend, we also have active reason to think that (i) is false. And if (i) is false then the total aggregate benefit of a single action is greater the total aggregate harm of that action, and so the total benefit of a million such actions will still be greater than the total harm of a million such actions.

Now, one can *conceive* of models on which there are each-we dilemmas even though any individual action does more good than harm. Perhaps, when considering a large number of actions, the goods are additive but the harms are multiplicative. But I see no reason why this should be so, I know of no empirical evidence suggesting that this is so, and Temkin gives us no reason to think this either.

CONCLUSION

Let me end with a comment about the nature of the broader dialectic regarding Singer's argument for the conclusion that we in rich countries have strong duties of beneficence. Often, critics of Peter Singer focus on whether or not aid is effective. But that is fundamentally failing to engage with core of Singer's argument. Correctly understood, that argument is about the ethics of buying luxury goods, not the ethics of global development. Even if it turned out that every single development program that we know of does more harm than good, that fact would *not* mean that we can buy a larger house, safe in the knowledge that we have no pressing moral obligations of beneficence upon us. There are thousands of pressing problems that call out for our attention and that we could make significant inroads on with our resources. Here is an incomplete list of what \$10,000 can do (noting, in each case, that any cost-effectiveness estimates are highly uncertain, with large error bars, and refer to expected value):¹⁰

- Spare 20 years' worth of unnecessary incarceration, while not reducing

10. For recommendations for individual donors for all these cause areas, see Open Philanthropy (2017a).

public safety, by donating to organisations working in criminal justice reform (Open Philanthropy Project 2017b).

- Spare 1.2 million hens from the cruelty of battery cages by donating to corporate cage-free campaigns (Open Philanthropy Project 2016).
- Reduce the chance of a civilisation-ending global pandemic by funding policy research and advocacy on biosecurity issues (Open Philanthropy Project 2014).¹¹
- Contribute to a more equitable international order by funding policy analysis and campaigning.¹²

In order to show that Singer's argument is not successful, one would need to show that for *none* of these problems can we make a significant difference at little moral cost to ourselves. This is a very high bar to meet. In a world of such suffering, of such multitudinous and variegated forms, often caused by the actions and policies of us in rich countries, it would be a shocking and highly suspicious conclusion if there were simply nothing that the richest 3% of the world's population could do with their resources in order to significantly make the world a better place.

The core of Singer's argument is the principle that, if it is in our power to prevent something very bad from happening, without thereby sacrificing anything morally significant, we ought, morally, to do so. We can. So we should.

REFERENCES

Barder, Owen. "Can aid work?: written testimony submitted to the House of Lords." *Washington, DC: Center for Global Development* (2011).

Burnside, Craig, and David Dollar. "Aid, policies, and growth." *American Economic Review* 90, no. 4 (2000): 847-868.

11. I don't provide a cost-effectiveness estimate here, but if we assume that \$10 billion of carefully targeted philanthropy would reduce extinction risk by 0.1 percentage points over the coming century, that would mean that \$10,000 gives a one in a billion chance of preventing an extinction-level event. That's an expected 7 present-day lives saved, and a prevention of the expected loss of life of five million future lives who otherwise wouldn't have been born; this latter estimate assumes constant population levels until the Earth is no longer habitable in 600 million years.

12. For example by donating to the Center for Global Development: see Open Philanthropy Project (2016b). Again I don't provide a cost-effectiveness estimate, but Open Philanthropy believe that over their lifespan CGD have caused billions of dollars worth of value for very poor countries on the basis of a \$150 million budget.

Clemens, Michael A., Steven Radelet, Rikhil R. Bhavnani, and Samuel Bazzi. "Counting chickens when they hatch: Timing and the effects of aid on growth." *The Economic Journal* 122, no. 561 (2012): 590-617.

Deaton, Angus. *The great escape: health, wealth, and the origins of inequality*. Princeton University Press, 2013.

Easterly, William. 'Some cite good news on aid.' *Aid Watch* (2009), accessible at <http://www.nyudri.org/aidwatcharchive/2009/02/some-cite-good-news-on-aid> [Accessed 22/5/19].

Gamble, Carol, Paul J. Ekwaru, Paul Garner, and Feiko O. Ter Kuile. "Insecticide-treated nets for the prevention of malaria in pregnancy: a systematic review of randomised controlled trials." *PLoS Medicine* 4, no. 3 (2007): e107.

Gettleman, Jeffrey. "Meant to keep malaria out, mosquito nets are used to haul fish in." *New York Times* 24 (2015).

Givewell. 'Giving cash versus giving bednets.' www.givewell.org (2012), accessible at <https://blog.givewell.org/2012/05/30/giving-cash-versus-giving-bednets/> [Accessed 22/5/19].

———. 'A Conversation with Abraham Mnzava on October 18, 2013.' www.givewell.org (2013), accessible at <https://www.givewell.org/files/conversations/Abraham%20Mnzava10-%202018-13%20%28public%29.pdf> [Accessed 22/5/19].

———. 'Putting the problem of bed nets used for fishing in perspective.' www.givewell.org (2015), accessible at <https://blog.givewell.org/2015/02/05/putting-the-problem-of-bed-nets-used-for-fishing-in-perspective/> [Accessed 22/5/19].

———. 'How GiveWell uses cost-effectiveness analyses.' www.givewell.org (2017), accessible at <https://blog.givewell.org/2017/06/01/how-givewell-uses-cost-effectiveness-analyses/> [Accessed 22/5/19].

———. 'Against Malaria Foundation.' www.givewell.org (2018a), accessible at <https://www.givewell.org/charities/amf> [Accessed 22/5/19].

———. 'Mass Distribution of Long-lasting Insecticide-Treated Nets.' www.givewell.org (2018b), accessible at <https://www.givewell.org/international/technical/programs/insecticide-treated-nets> [Accessed 22/5/19].

———. 'Counterfactual in countries where AMF funded distributions.' www.givewell.org (2018c), accessible at https://www.givewell.org/charities/against-malaria-foundation/supplementary-information#Counterfactual_in_countries_where_AMF_funded_distributions [Accessed 22/5/19].

———. 'Revisiting leverage.' www.givewell.org (2018d), accessible at <https://blog.givewell.org/2018/02/13/revisiting-leverage/> [Accessed 22/5/19].

Glennie, Jonathan, and Andy Sumner. "The \$138.5 billion question: when does foreign aid work (and when doesn't it)." *CGD Policy Paper* 49 (2014).

Global IDEA Scientific Advisory Committee. "Health and economic benefits of an accelerated program of research to combat global infectious diseases." *Canadian Medical Association Journal* 171, no. 10 (2004): 1203-1208.

Lengeler, Christian. "Insecticide-treated bed nets and curtains for preventing malaria." *The Cochrane Library* (2004).

MacAskill, William. *Doing good better: Effective altruism and a radical new way to make a difference*. Guardian Faber Publishing, 2015.

———. (forthcoming). "The definition of effective altruism," in Greaves and Pummer (eds) *Effective Altruism: Philosophical Issues*, OUP.

Mekasha, Tseday Jemaneh, and Finn Tarp. "Aid and growth: What meta-analysis reveals." *The Journal of Development Studies* 49, no. 4 (2013): 564-583.

Ng, Qin Xiang, Shawn Shao Hong Koh, Hwei Wuen Chan, and Collin Yih Xian Ho. "Clinical use of curcumin in depression: A meta-analysis." *Journal of the American Medical Directors Association* 18, no. 6 (2017): 503-508.

Open Philanthropy Project, 'Biosecurity,' (2014), accessible at

<https://www.openphilanthropy.org/research/cause-reports/biosecurity> [Accessed 22/5/19].

———. 'Effectiveness of corporate campaigns,' (2016a), accessible at https://www.openphilanthropy.org/focus/us-policy/farm-animal-welfare/humane-league-corporate-cage-free-campaigns#Effectiveness_of_corporate_campaigns [Accessed 22/5/19].

———. 'Center for Global Development — General Support 2016,' (2016b), accessible at <https://www.openphilanthropy.org/focus/global-health-and-development/miscellaneous/center-global-development-general-support-2016> [Accessed 22/5/19].

———. 'Suggestions for Individual Donors from Open Philanthropy Project Staff - 2017.' (2017a), accessible at <https://www.openphilanthropy.org/blog/suggestions-individual-donors-openphilanthropy-project-staff-2017> [Accessed 22/5/19].

———. 'Criminal Justice Reform Strategy: What we're doing and why.' (2017b), accessible at https://www.openphilanthropy.org/focus/us-policy/criminal-justice-reform/criminal-justice-reform-strategy#What_were_doing_and_why [Accessed 22/5/19].

Protopopoff, Natacha, Jacklin F. Moshia, Eliud Lukole, Jacques D. Charlwood, Alexandra Wright, Charles D. Mwalimu, Alphaxard Manjurano et al. "Effectiveness of a long-lasting piperonyl butoxide-treated insecticidal net and indoor residual spray interventions, separately and together, against malaria transmitted by pyrethroid-resistant mosquitoes: a cluster, randomised controlled, two-by-two factorial design trial." *The Lancet* (2018): 1577-1588.

Rajan, Raghuram G., and Arvind Subramanian. "Aid and growth: What does the cross-country evidence really show?." *The Review of Economics and Statistics* 90, no. 4 (2008): 643-665.

Ravallion, Martin. "On the role of aid in The Great Escape." *Review of Income and Wealth* 60, no. 4 (2014): 967-984.

Roodman, David. "The impact of life-saving interventions on fertility," <https://davidroodman.com> (2014), accessible at <https://davidroodman.com/blog/2014/04/16/the-mortality-fertility-link/> [Accessed 22/5/19].

Tiono, Alfred B., Alphonse Ouédraogo, Daouda Ouattara, Edith C. Bougouma, Sam Coulibaly, Amidou Diarra, Brian Faragher et al. "Efficacy of Olyset Duo, a bednet containing pyriproxyfen and permethrin, versus a permethrin-only net against clinical malaria in an area with highly pyrethroid-resistant vectors in rural Burkina Faso: a cluster-randomised controlled trial." *The Lancet* (2018).