



---

## William MacAskill's interpretation of AMF cost-effectiveness estimates in Doing Good Better

8 messages

---

**Alexey Guzey** <alexey@guzey.com>  
To: info@givewell.org

Sun, Nov 18, 2018 at 10:45 PM

Hi,

I recently wrote a [post](#) in which I point out various misquotes, misrepresentations, etc in William MacAskill's *Doing Good Better*.

In the post I point out that in the comments section of [Why We Can't Take Expected Value Estimates Literally \(Even When They're Unbiased\)](#), Holden Karnofsky writes:

There are several cases in which Giving What We Can appears to take cost-effectiveness estimates literally even though they do not include adjustments of the sort described here. The quote from Will in the above post is one example; another is the [statement](#) that "Charities which focus on STH, schistosomiasis, and LF are at the very top end of cost-effectiveness — about 100 times more effective than typical developing-world health interventions."

Here, he writes that writing about "100 times" the effect seems to require cost-effectiveness estimates literally.

In DGB, MacAskill writes:

the QALY allows us to make comparisons across very different programs that combat very different illnesses. By donating to the Against Malaria Foundation, which buys and distributes long-lasting insecticide-treated bed nets, you would, by this estimate, provide **five hundred times the benefit** as you would by spending the same amount of money treating Kaposi's sarcoma. **[emphasis mine]**

MacAskill uses [GiveWell's AMF cost-effectiveness estimate](#) here. It seems that he takes this cost-effectiveness literally, by writing about "five hundred times the benefit", as per Karnofsky's interpretation of GWWC's "100 times more effective" statement about deworming effectiveness.

As of DGB's writing, GiveWell's AMF page said:

As a general note on the limitations to this kind of cost-effectiveness analysis, we believe that [cost-effectiveness estimates such as these should not be taken literally](#), due to the significant uncertainty around them.

Here, GW specifically asks not to interpret the AMF estimate literally, which MacAskill does in the book.

On the basis of these three quotes, I concluded that MacAskill misrepresents GW's estimate.

In a [response](#) to my post on the Effective Altruism Forum, MacAskill writes:

I asked an employee of GiveWell to look over the relevant parts of the manuscript of DGB before it was published; in the end five employees did so, and they were happy with how I presented GiveWell's views and research.

Which to me seems inconsistent with the quotes from MacAskill, GiveWell, and Karnofsky that I provided. Is my interpretation of anybody's words or positions incorrect here?

Best,  
Alexey

---

**Alexey Guzey** <alexey@guzey.com>  
To: info@givewell.org

Wed, Nov 21, 2018 at 11:02 AM

Hi again,

Do you have any updates on this issue?

[Quoted text hidden]

---

**Catherine Hollander** <catherine@givewell.org>  
To: alexey@guzey.com

Thu, Nov 22, 2018 at 4:06 AM

Hi Alexey,

Thanks for reaching out to GiveWell; apologies for the delay in response.

We're unfortunately quite capacity constrained this week. We are planning to respond soon.

Thanks for your patience.

Best,  
Catherine

[Quoted text hidden]

—

Catherine Hollander

[www.GiveWell.org](http://www.GiveWell.org)

---

Alexey Guzey <alexey@guzey.com>  
To: catherine@givewell.org

Thu, Nov 22, 2018 at 11:05 AM

Thanks!  
[Quoted text hidden]

---

Alexey Guzey <alexey@guzey.com>  
To: catherine@givewell.org

Fri, Nov 30, 2018 at 8:17 PM

Hey Catherine,

Do you have any updates on this?  
[Quoted text hidden]

---

James Snowden <james.snowden@givewell.org>  
To: alexey@guzey.com

Mon, Dec 3, 2018 at 1:19 PM

Hi Alexey,

Thanks for getting in touch. I'm afraid I don't have the capacity to respond to everything in your post, but I'll respond to the points you raise in this email. I also can't commit to continuing any future correspondence.

For full transparency, I worked for Giving What We Can/the Centre for Effective Altruism between 2015 and 2017. I was not involved in the production of *Doing Good Better*. Since 2017, I've been working as a research consultant for GiveWell.

I think the most helpful way I can respond is by answering two questions:

1. What do we mean by "not taking cost-effectiveness estimates literally"?
2. On the basis of the quotes you sent, to what extent (if at all) did Will misrepresent GiveWell?

#### **What do we mean by "not taking cost-effectiveness estimates literally"?**

GiveWell's mission is to direct charitable funding to where it can help people the most, according to our criteria, while maintaining very high levels of transparency and honesty in our communication.

In an abstract sense, the thing we care about most is cost-effectiveness (distinct from *explicit cost-effectiveness estimates*). As it's infeasible to explicitly model every relevant consideration, we don't believe our cost-effectiveness estimates should be taken to represent our *true best guess* of a program's impact, and we take other considerations (e.g. our subjective interpretation of an organization's "quality") into account when making decisions (for impact-based reasons). Having said that, our cost-effectiveness estimates are one of the most important tools we use in our decision-making.

Since Holden wrote the blog post you reference in 2008, we've started to explicitly take more factors into account in our cost-effectiveness estimates, and place correspondingly more weight in those estimates (although we don't have an explicit formula for precisely *how much* weight). I believe this makes our estimates fairly distinct from most published estimates in the academic literature which (perhaps understandably, given the need to standardize assumptions for comparability) don't attempt to incorporate more subjective judgments calls.

We still qualitatively stand by the overall point made in that post, but recognize when to explicitly quantify, and when to apply looser heuristics is a difficult (and unsettled) problem of decisionmaking. We also recognize that when to do so will depend on (a) the particular consideration you're attempting to weigh; and (b) the research capacity of your organization.

Some examples of the increasing number of factors our cost-effectiveness estimates take into account:

1. We apply (often fairly subjective) "replicability" and "external validity" adjustments to our cost-effectiveness analysis, to account for our view of whether a study would replicate in the same, or different, contexts. We don't always believe these adjustments are "objectively justified" to the same extent as some of the inputs into our analysis, but taking our best informed quantitative guess forces us to think through and weigh as many of the relevant considerations as possible. As an example, see our [writeup](#) on GiveWell standout charity Zusha!, and the justification for the internal validity adjustment [here](#).
2. We now explicitly attempt to incorporate the possibility of crowding in and crowding out other funders from our programs. See [this post](#) for more details. In that post, we describe the tradeoff between explicit quantification and more intuitive decisionmaking:

"First, these estimates rely on more guesswork than most of our cost-effectiveness analysis, reflecting a fundamental tradeoff we face in deciding which considerations to explicitly quantify. Quantification forces us to think through not just whether a particular consideration matters, but how much it matters relative to other factors, and to be explicit about that. On the other hand, incorporating very uncertain factors into our analysis can reduce its reliability, give a false impression of certainty, and make it difficult for others to engage with our work. In this case, we thought the benefits of explicit quantification outweighed the costs."

We still don't attempt to create explicit quantitative models of *all* the considerations that play a role in our decisionmaking, and don't expect to do so in the future (although we'll often use rough back of the envelope calculations to help calibrate ourselves).

We've found explicit quantification tends to be less helpful when the consequences of a particular observation aren't well-defined. For example, our subjective assessment of organizational quality is a proxy for how well those organizations make a variety of decisions across a range of domains, which won't be captured in our cost-effectiveness analysis. The stakes of these decisions are unknown *ex ante*, which limits the usefulness and feasibility of a fully explicit analysis. In these cases, we think it's more honest to state our real thought process for a decision (e.g. on balance, we think the small difference in estimated cost-effectiveness is outweighed by the higher quality of this organization), rather than attempt to justify them through a quantitative model.

In sum, we don't think our cost-effectiveness estimates should be taken as our *true best guess*. But they do incorporate a large number of factors relevant to our decisionmaking, and are intended to reflect a reasonable approximation of reality. Depending on your interpretation of the word "literal," this may or may not imply that they should be taken "literally" (we had the former definition in mind).

### On the basis of the quotes you sent, to what extent (if at all) did Will misrepresent GiveWell?

*To give some context on GiveWell's thinking behind how we present our cost-effectiveness estimates:*

We've deemphasized our precise cost-effectiveness estimates on our website because unfortunately (but in my opinion understandably and non-maliciously) people tend to misinterpret concrete numbers as conveying more certainty than we think is warranted.

*To focus on the quotes you provide:*

>"the QALY allows us to make comparisons across very different programs that combat very different illnesses. By donating to the Against Malaria Foundation, which buys and distributes long-lasting insecticide-treated bed nets, you would, by this estimate, provide five hundred times the benefit as you would by spending the same amount of money treating Kaposi's sarcoma."

Here, Will explicitly says "by this estimate". I think that's a reasonable caveat in a popular book and goes a long way to differentiate Will's statement from the earlier quote you provide. I personally would have put more explanation of the uncertainty in a footnote, including noting that I hadn't personally vetted the analyses in question (if that was the case). But I've rarely seen a caveat like that outside GiveWell's website, and wouldn't generally expect it.

>"I asked an employee of GiveWell to look over the relevant parts of the manuscript of DGB before it was published; in the end five employees did so, and they were happy with how I presented GiveWell's views and research."

This is true.

In sum, Will's presentation is not exactly how I (or a typical GiveWell researcher) would have presented the comparison, but I think the way he presented it seems reasonable in a popular book aimed at a general audience. As I'm sure you understand, it's difficult to communicate about these complex topics effectively without losing the nuance. I do not believe Will misrepresented GiveWell.

All the best,

James

On Thu, Nov 29, 2018 at 5:27 PM Catherine Hollander <[catherine@givewell.org](mailto:catherine@givewell.org)> wrote:

[Quoted text hidden]

--

Catherine Hollander

[www.GiveWell.org](http://www.GiveWell.org)

--

James Snowden  
Consultant Researcher  
[www.GiveWell.org](http://www.GiveWell.org)

---

**Alexey Guzey** <[alexey@guzey.com](mailto:alexey@guzey.com)>

To: [james.snowden@givewell.org](mailto:james.snowden@givewell.org)

Mon, Dec 3, 2018 at 3:34 PM

Hi James,

Thank you so much for your response and for the explanation of GiveWell's position on the interpretation of cost-effectiveness estimates. Would it be okay for me to post this email correspondence in public and to comment on it? I think it would be a useful addition to my post.

It appears that I will need to retract my point on William's interpretation of GW estimates / maybe will need to retract the entire post.

[Quoted text hidden]

**James Snowden** <james.snowden@givewell.org>  
To: alexey@guzey.com

Mon, Dec 3, 2018 at 3:50 PM

Thanks Alexey,

I'm happy for you to share this email publicly and comment on it. I appreciate your transparency!

All the best,

**James**

[Quoted text hidden]