

## GiveWell San Francisco Research Event December 13, 2016 – Top Charities

*This transcript was compiled by an outside contractor, and GiveWell did not review it in full before publishing, so it is possible that parts of the audio were inaccurately transcribed. If you have questions about any part of this transcript, please review the original audio recording that was posted along with these notes.*

00:00 Elie Hassenfeld: Thanks for coming. I'm Elie Hassenfeld. I'm one of GiveWell's co-founders. Just want to give you a brief overview of how tonight's going to go. So I'm going to talk about GiveWell's work for the next 40 minutes or so. Then Holden, who is GiveWell's other co-founder is going to talk about our work on the Open Philanthropy Project. And then we're going to break off and a bunch of GiveWell and Open Philanthropy staff are going to stick around and you can see that we put signs on each of these conference room doors with the people who'll be in the rooms and what they know about, is there just to give you an opportunity to ask folks more questions one-on-one about whatever is on your mind.

00:48 EH: Just FYI, we're recording this, that's something we do with most of these sorts of things because a major goal of ours is to be as transparent as possible about the work we're doing. If you say anything in this meeting and you'd rather we remove it from the recording, we're really happy to do that so just either let me know, after the event, let Catherine know, she's right over there or you can just email [info@givewell.org](mailto:info@givewell.org) after the fact, and we can just remove whatever you'd like us to remove from the recording.

01:19 EH: We put a lot of information on our website and so our goal in hosting these events is to give you a chance to ask questions about the things that are on your mind that you're not able to find on the website or you've just been wondering about. So really feel free to ask whatever you like, whatever you are wondering about even if you think it's something that we probably wrote about somewhere already. We're also happy, I'm very happy to use this time to just try to help better communicate what we do and why we do it.

01:50 EH: I am planning to target this more at a group of people who are relatively familiar with GiveWell already because looking around the room, I think that describes a fair number of you who are here. If you find yourself really lost or have basic questions, I think talking to staff one on one is a really good way to get up to speed. And also, we appreciate feedback after the fact, and so we're going to send out a survey to folks who attended and we just appreciate hearing from you about whether this was targeted at the right level or if there's anything that you prefer we do differently. So, the main thing that I want to cover, I just want to quickly give an update on the recommendations we recently released and our updated top charities. I want to talk a bit about the relative strengths and weaknesses of these top charities and what sorts of considerations I think donors should have in mind if they're choosing between them and then I want to just very briefly touch on some of the future work I think we're going to do.

02:52 EH: So, briefly our mission is to find and recommend outstanding charities. Each December, in advance of the time of year when so many people give so much to charity, we release updated list of recommendations which is the short list of the groups that we most recommend, given GiveWell's criteria. And so, in this past year, we recommended again the four organizations that we have recommended in past years. Two which work on deworming programs, these are programs that treat primarily children in the developing world for parasitic infections.

03:31 EH: So these groups are the Schistosomiasis Control Initiative and the Deworm the World Initiative. We recommended a group called, the Against Malaria Foundation again, they support malaria net distributions in Africa. And then we also had recommended and continue to recommend a group called GiveDirectly, and they just give cash directly to the poorest people in the world. In addition of those four, we added three new organizations this year, two which also work on deworming, the END Fund and Sightsavers, and then a group that works on an intervention we hadn't recommended before. So this is a group called Malaria Consortium and they run a program called, Seasonal Malaria Chemoprevention. And this is distributing pills to children during this time of year when malaria is most common to prevent cases of malaria.

04:23 EH: And this intervention seems, roughly similarly, cost-effective to a malaria net program. So probably initially, you noticed four of these top charities are deworming organizations, and I think that just very briefly I want to explain to some extent why that is and how that relates to the process that we follow. So we're really assessing organizations on three primary criteria: Evidence of effectiveness, how do we know that the program itself has a causal impact and improves people's lives; cost-effectiveness, what's our best guess about a quantitative estimate of the cost to improve a life by a certain amount or save a life; and then finally room for more funding, the question of how and how effectively would organizations utilize the additional funds that we expect them to receive on the basis of GiveWell's recommendation.

05:17 EH: And we start that process by looking at academic literature, literature that's totally independent of any individual charitable organization looking for evidence that the interventions we're looking for or this better program is effective and cost-effective. And the reason we start there is that most charitable organizations will not have the sort of information available to them that could directly demonstrate the causal impact of the program that they're running. After that, we'll look for charities that are implementing those programs. And so, in the case of deworming, we think the independent literature collectively makes a very strong case that deworming is a very good value for money, and so we then are looking for organizations that can utilize funds to scale-up deworming programs.

06:10 EH: And then when we find organizations running those programs, we then look for evidence that the organizations can effectively implement those programs, that they're doing so in a way that we think would lead to strong effects, and that they have the capacity to utilize additional funding. That's the rough process that we've been following throughout GiveWell's history, which is now about, we're about 10 years old, and we have these seven recommendations that we just released. So in terms of the sort of relative trade-offs between these organizations, on one hand, there's GiveDirectly. And I think for donors, what they do is just give cash directly to very poor people. So I think for donors who either say, "I'm particularly skeptical about charities' abilities to have impact and I really want to find something that I'm highly confident has significant impact," or for donors who I think are understandably nervous about trying to send money thousands of miles overseas and have a large effect, GiveDirectly which just puts money in the hands of some of the poorest people in the world and lets them decide how to give, offers a really good option. Because it's arguably the simplest most straight forward way to give charity.

07:35 EH: In addition of all the organizations that I've interacted with or considered over the time that GiveWell's existed, as an organization, GiveDirectly really stands out. I think they are by far just exceptional as an organization in terms of their communication, their monitoring, their looking for problems that they try to fix, their ability to scale and find new ways to implement their programs effectively. One nice fact about GiveDirectly is that, for all the money they give out as part of their cash transfer program, they're also including that money as part of research experiments, where they're trying to learn more about how to give aid effectively. So one type of thing they might do is run an experiment or have their transfers participate in an experiment where researchers are trying to compare the effectiveness of a direct cash transfer to some other sort of development program like a training program or a livestock delivery program.

08:41 EH: And so this is the sort of attitude that GiveDirectly has as an organization, and I think they're really exceptional for it. The downside is that we think that ultimately cash seems to us to be less of cost-effective on a per dollar basis than the other programs we recommend. And it's worth saying that the cost-effectiveness estimates that we try to reach are subject to a large number of debatable assumptions where the data is quite poor meaning, what is the baseline income of the people that GiveDirectly is serving. We have some information about that but it's not particularly high-quality data. And this type of issue runs through all the organizations we consider. These things are also subject to some very clearly subjective philosophical value judgments. Like how do you value the benefit of saving a child's life from malaria? Which is what the Against Malaria Foundation accomplishes versus the value that GiveDirectly accomplishes which is increasing someone's available cash over some period of time.

09:56 EH: And so in our cost-effectiveness model we try to put explicit values on these trade-offs, it's on our website, you can go through it if you'd like. But we ultimately come to the conclusion that GiveDirectly is less cost-effective than these other options, but I think it's a very reasonable position to hold that these cost-effectiveness estimates are the output of a somewhat difficult to create debatable spreadsheet. And for donors who really find that to be not a very compelling form of analysis, I think GiveDirectly also is a particularly good option. So then there are the deworming groups, and I think what makes deworming stand out is it is the intervention where I think we have the lowest overall confidence that it has a significant impact. But the impact that it might have, if it is successful seems so large and the intervention is so cheap that we recommend it.

10:55 EH: And so the possible impact that deworming could have and this comes largely from a single experiment run in Western Kenya in the late 1990s, is that deworming children could lead them to have something in the range of a 25% increase in income as adults. And so the idea is that maybe this program causes some sort of cognitive or physical development by removing the parasite that causes this large long-term effect. But this result largely comes from one study, and other short-term studies don't find noticeable significant short-term effects. And so we have a significant question about the mechanism through which deworming might work. But nonetheless it's extremely cheap in the range of 50 cents to \$1 per child per year, you might have this, frankly, almost incredible impact long term, and so in our model we try to build in a fair number of doubts and adjustments for the fact that we're uncertain. But if you're not particularly worried about the risk of limited impact and you're maybe more looking at in an expected value framework, "Where can I think I'll get the most bang for our buck?" I think deworming offers a pretty good option on that front.

12:22 EH: And then finally we have the two groups focused on malaria prevention, The Against Malaria Foundation which distributes malaria nets and Malaria Consortium, which runs Seasonal Malaria Chemoprevention. Where these stand out, I think both of these organizations are somewhere in between the case for cash and the case for deworming, in terms of the likelihood that they have impact and sort of their aggregate cost-effectiveness according to our model, but these are the only life-saving interventions that we're recommending. We think cash gives people the capabilities to buy things that they need. Deworming, the effect that we're really focused on is the potential that it increases income later in life, and the malaria-focused interventions are averting the deaths of young children. And so, I think for people who put very high value on life-saving interventions relative to life-improving or wealth-increasing interventions, I think the malaria groups win.

13:25 EH: For anyone who's really interested in diving into all the details, the cost-effectiveness model to kind of play around with all of these assumptions is on our website. And then finally, for people who just want to follow our recommendation, our current recommendation is to give 75% to the Against Malaria Foundation and 25% to SCI, and that allocation is a function of the amount of funding that we expect individuals to give due to our recommendation over the next few months, and all seven organizations, current funding levels and our understanding about what their needs are. And so, that's our sort of bottom-line recommendation for people who just want to follow our current advice.

14:09 EH: Very briefly going to touch on our future and then we'll move into questions. So, the main thing I want to talk about is just what we're doing to find new GiveWell top charities. And in addition to continuing our standard process, we're also working on an area of work, we've called today our sort of experimental process, and the idea of this program is to try to support the development of future top charities. Because GiveWell itself recommends organizations that are already at a certain stage in their lifecycle. They've already built up a track record. There's already a research base. They've already done the monitoring that demonstrates that they're effective and that there are a whole host of organizations that could potentially reach that level, but haven't yet gotten there.

14:57 EH: And so, we'd like to provide some support to enable them to get there. And I really just want to focus... There's a whole bunch of these and I'm happy to talk more about it, but one of the opportunities I'm really excited about is work we're doing with a group called IDinsight. And Buddy Shah is here. He's the CEO of IDinsight. And they're a group that we see potentially as a very useful and effective partner in helping us on this work. So IDinsight and Buddy, forgive me if I don't describe this well, it's just my way of doing it, is a group that focuses on organization in context-specific evaluations with the direct aim of influencing policy makers and funders. And so, GiveWell and its donor community is a group of funders or donors that IDinsight is really interested in trying to influence. And so, we've just started a project with them where there's a small team of IDinsight staff that's working directly with us to try and improve the information that's available about potential GiveWell top charities. So two quick examples of how this might work. One area of work that we've been really interested over time is performing surgeries to improve quality of life.

16:17 EH: And one type of surgery that is very cheap and very effective is cataract surgery, where it's inexpensive, but when we've talked to the organizations that implement this program, they have had very limited monitoring of how well their program works. Does their program actually lead to more surgeries being performed or would people have access to the surgeries even independent of that organization existing? What are the pre and post-operative outcomes or the pre-operative condition and the post-operative outcome for people who are being served by the organization? And so, one sort of project we've discussed working with IDinsight on is getting better data both pre-surgery and post-surgery on how well the organizations do in running their programs.

17:13 EH: And then another area is a group called New Incentives, which is now running a program where they provide conditional cash transfers to parents with young children to incentivize them to come and have their children receive childhood immunizations. And there's very good evidence that if you receive an immunization, it improves your life and it prevents the potential for a childhood disease. But what we don't know is whether new incentives providing these incentives will cause immunization rates to rise. And so, we've been talking to IDinsight to see whether we can work with them to design an experiment where we would directly measure the effect that New Incentives program has on immunization rates and, depending on the results of that study, we'd be in a position to decide whether or not to recommend them as a top charity. So I'm going to pause there and open the floor up for questions. Yep?

18:15 Speaker 2: Is GiveWell or any its affiliated organizations investing money into reducing the uncertainty around the impact of deworming?

18:23 EH: So the question is, "Is GiveWell investing in reducing uncertainty around deworming?" We're talking about doing... We're considering one thing, and then there's another thing which I'll tell you about, and then we've talked about another thing which I think is not the best opportunity for us. So the way in which we have the results I described for deworming, are these long-term follow-up surveys from this initial randomized controlled trial that happened in Kenya in the late '90s. And the researchers have gone back and resurveyed the population that was served by this program, every several years, to get updated data.

19:03 EH: In the most recent survey they only went back to something like 20% of the people who were part of the program. We've been talking to the researchers about whether we could provide some resources so they could survey everyone. Because we'd rather see the full sample as opposed to just a partial sample. And this is very attractive to us because getting the data only requires additional funding, as opposed to starting another experiment from scratch. We've also thought about whether we should try to support another deworming randomized controlled trial that could be a second experiment that we could look to years later about whether deworming has the effects that we think it might.

19:46 EH: And the problem there is, first, kids are dewormed in childhood and then start to earn wages 10 years later. So the lag time is extremely long. And I think that one risk is that by the time you got the results, the world would have improved enough in terms of how much funding would have gone to deworming regardless, that it wouldn't even be worth waiting around for that data. The second issue is that the study that we would need to run, would be significantly larger than the study that was run in Kenya in 1998 and 1999. And that's because, in that study, worm rates were extremely high, relative to everything that exists today. Both because, I think it was so long ago that worm rates were much higher, and it happened in an El Nino year, and so there was a great deal of flooding. And so worm rates actually spiked in the second year of the study, relative to the first year.

20:47 EH: So in order to run a study that could have the statistical power to pick up a similar effect, given today's levels of worm rates, would be prohibitively expensive. And it would require a sample far larger than what they had done initially. And so for that reason, we haven't tried to run additional trials on deworming. The final thing we've thought about, but haven't made much progress on, is whether there's some other intermediate indicator you could test for, that might potentially give us useful information. And that's something we talked a bit to IDinsight about. It's still something that's in the back of our minds, but we haven't found the right opportunity to pursue. Yep?

21:30 Speaker 3: So, I have been actually following GiveWell for a few years, and I feel your recommendations are always pretty consistent. And I just wonder what are your thoughts for other organizations that are a bit different from the ones you recommend. Specific examples are maybe organizations that have improved the survivability for people who... Immigrants from Africa to Europe [21:55] \_\_\_\_\_. And also organizations that provide shelters and basic healthcare in addition to people from poor countries. Because I didn't see too much of this type of organizations on your website. So I just wonder, whether you feel they are less cost-effective, or what are some things to look for beyond the ones you do recommend.

22:19 EH: And I'm just repeating these questions for the recording. The question is, "GiveWell's recommendations are clearly of a particular type. What about other organizations?" GiveWell's goal is to maximize the impact of the recommendations we make, subject to the criteria of strong evidence in particular, and then the others, cost-effectiveness and room for more funding. We're explicitly ready to potentially leave some impact on the table for the sake of having recommendations that are easily vettable and verifiable by ourselves internally and the outside world. There's certainly, kind of conceptually, the idea that you could have opportunities that don't meet those criteria but could be very high impact. It's something that everyone at GiveWell agrees with.

23:18 EH: And that notion is a lot of the basis for the work we're doing on the Open Philanthropy Project, which Holden will talk about next. The main thing I think is worth noting, is that when you give up verifiability and strong evidence, you are giving up a lot in terms of your ability to know whether or not something worked. And when we give that up, what we then really want to look for, are the projects that potentially maximize impact, as to the greatest extent possible. And so looking for the, just as an example, the scientific research program that could have the biggest long-term effect on the world's ultimate, how well it goes. And so then the main thing we think about is that we like to be at both ends of the spectrum. Either at the highly vettable, verifiable end or just trying to maximize the expected value, even if that means that the risk of any particular project failing is very high.

24:27 Speaker 4: Pardon me if this is already on the website but I was wondering, do you consider externalities to interventions? So for example, if there's a malaria net intervention, what does that do to local people who might produce malaria nets? Or also dependence on this aid in the form of this malaria nets going forward for these people?

24:51 EH: Yeah. So the question is about externalities, and maybe negative or offsetting effects of our top charities. And just again, if you think we might just ask anything even if it's on the website, that's in the fourth bullet of a 300 footnote page and so it's hard to find. Yeah, so it's something we think about though I don't think we have the easiest, the clearest most direct way to demonstrate that it's not a problem. And I think that it's clear that in this case, on some level, when external donors provide bednets you reduce the incentive that locals would have to create bednet companies, you make it less likely that local governments will provide bednets on their own.

25:38 EH: In the case of our top charities, I think two things are true. While those effects seem true to me, that they have that sort of negative effect, at the same time I don't think there's much evidence or any evidence that local countries are able or willing to pay for malaria nets. So for example, we've looked at cases where the Against Malaria Foundation considered providing funding to a country, and then for whatever reason didn't ultimately provide funding. Did that country then come through and say, "We'll provide funding now"? And in the local countries are just not stepping forward. Obviously, that itself is not proof because there's a larger ecosystem around it and they have expectations, but I don't think there is... I think the available evidence looks like, the countries are not able to provide this sort of funding on their own.

26:34 EH: And the same is true for the local bednet manufacturers. The goal that the Against Malaria Foundation has is to purchase these nets at the lowest available price. And when there have been African bednet manufacturers that have propped up and people have said... Some donors will say, "You have to use the African bednet manufacturer," because we should purchase it locally, they end up costing more. I think the argument I would make is that in this case, the short-term benefits seem clear enough and large enough that it sort of outweighs what we know about the potential negative effect, but I don't think we have the most definitive answer to that.

27:23 Speaker 5: If evidence shows that a particular program is effective where it was initially tested, but additional funding would induce that organization to try to replicate that program elsewhere but it's unknown whether the context in that next place would allow it to be as effective, how do you think about that? What do you do with that one?

27:41 EH: Yeah. So the question is, let's say the evidence is focused on a program in one area but now the charity is working in another area. One thing is that the charities we recommend primarily are working on health programs. And so the external validity of health programs, I think, is intuitively going to be higher than a more complicated, or intervention that relies more on the cultural context. So a pill that prevents malaria is likely to work the same way in Nigeria as it is in Tanzania.

28:17 EH: That said, there certainly are ways in which, even with our top charities, the context in which they're working differs substantially from the place where the evidence was first gathered. I think deworming is the best example. I said worm rates are really high. It's far higher than anywhere else. You can easily imagine a model of the impact that the program has where at worm rate levels that are 10% of the initial program, just applying a 10% adjustment factor doesn't sufficiently take into account that difference. There are also a number of worms, and so in the study in Africa, there was a particularly one sort of worm infection, Schistosomiasis was present, and in India where the Deworm the World Initiative works, there is no Schistosomiasis. And so we do try to adjust for this but I think it's a fair point in a way in which the programs implemented by our charities don't, by any means, perfectly match the evidence that was available. And so we try to do our best to match it up as best we can to what the initial evidence was.

29:27 Speaker 6: Can you just share how you consider the interventions that you're going to research? I'd heard that you were interested in Vitamin A supplementation. Is that still something you are looking for organizations that will roll it out? And are there other interventions that you're thinking about?

29:47 EH: So when you say "how do you consider", do you mean where does the list come from, or how do we decide which ones to look into?

29:54 S6: Exactly.

29:55 EH: Okay. So we have a very long list of programs that we're considering doing additional research on, and we try to gather that list from the sort of cast the net as wide as we possibly can to anything you could imagine, from major publications that list interventions to anything that charities tell us they do to just following ongoing papers that are released in public health, global health journals. And then we have that list and we try to use in that list just prioritize the looking into those interventions based on our best guess at how good a fit they'll ultimately be for a GiveWell

recommendation in terms of the strength of the evidence and the cost-effectiveness. So Vitamin A is an intervention that we looked at, we have a report on, and for that one, in particular, we definitely are still actively looking for organizations that are implementing it. Yep?

30:51 Speaker 7: Is anything changing in the second-tier charity list? Any surprises or interesting things that turn up?

30:56 EH: So we added two additional groups there both that work on micronutrient fortification but I don't think there's any really major news with the second-tier list so to speak. What we're trying to do with that list is say, "Here are organizations that run programs that have strong evidence or stand out in some sort of other special way but we haven't been able to see the full case that would make them competitive in our view with our top charities, but we still want to give them that recognition that of the 1000 organizations we've considered in some way during GiveWell's time, these additional ones do seem quite special to us." So they're not a group that we spend a great deal of time on but they're on our website. Yep, right there?

31:52 Speaker 8: Yep. I have a question. To what extent does GiveWell's sort of track or follow or consider the slightly downstream effects of the interventions that you promote where it might not be quite as numerically calculable what that effect is but it could have more of a qualitative than a quantitative effect? So an example might be with preventing malaria that then that reduces child deaths so it reduces family size, so that improves lives of women, so there's a kind of butterfly effect from that, in terms of slightly more personal or cultural things rather than just the number of lives saved per se. Is that something you look into or think about when you're choosing the interventions?

32:37 EH: Yeah. So the question is, how do we consider these downstream effects of our recommendations? I think it's not something that we explicitly try to quantify, and when we're trying to do the quantification, we're always balancing the ability to come up with an estimate that's at all reasonable and not just total guesswork with the information that we have. And I think as we move further and further downstream, it becomes harder and harder to have any sense of what that means. I also just intuitively guess that the first approximation that we have in our cost-effectiveness models are reasonably sort of an equivalent comparison across organizations. And I don't have any good reason to believe that the downstream effects from a cash transfer intervention that provided better nutrition to some children in the field would have a sort of a fundamentally different sort of effect than relative to the immediate impact, than something like malaria would. Okay. Yep?

33:46 Speaker 9: So I have an additional question. Could you give us an overview about how large this academic community is for providing this research for charities, the effectiveness of different kind of charities or different kind of treatment they can do to people because that's what your resources for evaluating the impact of different projects we've invested with. And in addition to the academic research, do you trust any other third-party independent research results?

34:19 EH: Got it. So the question is, what is sort of the nature of the academic community's research we rely on? The community is quite large. It's not really focused on, I would say, charitable organization specifically as much as international development more broadly. And many universities have people who are focused on global health, public health, development economics. In journals like The Lancet, New England Journal of Medicine, those are the sorts of places that are publishing research that we rely heavily on. And it's not necessarily just academic. There've been cases where charities themselves or think tanks or something are producing research that's useful, but by and large, the sort of academic research article on a program is far and away the sort of input that we rely most heavily on. Alright, Buddy.

35:17 Speaker 10: How if at all do you guys think about influencing other donors and influential spenders in international development to talk more about GiveWell world view and giving them mechanisms [35:29] \_\_\_\_.

35:30 EH: Yeah. So sort of what's our role in talking to say, the major government donors. It's not something we've done any of really to date, except when they have sought to talk to us and they say like, "Can you share how GiveWell thinks about this sort of problem?" It is something we've thought more about now for our future. I think that GiveWell for its first 10 years was primarily focused solely on research, and we weren't trying very hard to bring in more donors, and we were very fortunate with the media got, with the donors who found us, to get a fair amount of traction.

36:11 EH: And as we look towards our future, I think that the government aid agencies in some form collectively account for so much aid that I hope that we're able to influence them, either in sort of a direct way to support the organizations we recommend. And this is not crazy, the UK government gives funds about half of SCI, one of the deworming recommendations budget, totally independent of us but it's the sort of organization that's on their radar, or more generally just trying to influence them to take a more evidence-based approach to the programs that they support.

So check back next year and hopefully I'll have a better answer.

37:27 EH: Alright, it's just about 7:30, so I think we should pause here, take a brief break and then reconvene for Holden. I just want to say, I really appreciate that everyone who comes out here, this is sort of our version of the charity, I shouldn't say charity gala because it can't be our version of the charity gala, but for folks who are happy to sit around and debate charity as opposed to have fancy food and wear tuxedos, and we're really grateful for all the interest and desire to just engage in the substance of what we do. So thank you.