

GiveWell San Francisco Research Event, December 15, 2014 - 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.

[background conversation]

01:15 Elie Hassenfeld,: Alright. Hey, everyone, thanks for coming. We're about to get started. So, just a quick reminder to silence your phones. The basic way this event is going to work is we're going to put the first hour focus mostly on our traditional work, focused internationally top charities, and the second hour on some of the newer work we're doing on the Open Philanthropy Project. The way the first hour is going to work is I'll talk for a few minutes about some of our new top charities or some of our top charities and pause for questions, and then I'll talk about some of the other organizations we've recommended this year as standouts and pause for questions again.

01:54 EH: The reason we hold these events, or one of the main reasons, is that we know there's a lot of dense information on our website, and we want to create a way for people to engage who have not read all of the material on our website. So, if you have questions, feel free to ask them. Don't hold back just because you think it's something that you probably could find in a footnote on one of our charity reviews. We're really hoping that this is less of us talking and more of you asking questions and engaging with the things that we're saying.

02:24 EH: We're recording this event, consistent with our general practice of transparency. We're going to post a recording and a transcript on our website. If there's anything you say that you'd prefer for us to remove from either, we're happy to do that, so just let me know after the event. I'm Elie Hassenfeld, I'm one of GiveWell's co-founders. And this is Holden Karnofsky, the other co-founder. And so, I'll be talking about the traditional stuff, and then Holden will talk about Open Philanthropy.

02:56 EH: So, first I'm going to talk about our top charities. There's four this year and a lot of them are familiar names, but I'm going to go through them all in turn. There's basically major updates on all of them. And just as a reminder, the way that our traditional work goes, is that we first look for programs where charities that have very strong evidence of effectiveness, meaning we know they have an impact when they're implemented correctly. Then, we look for charities that can show us that they're accomplishing that. And we look for programs that are implementing, that are accomplishing a significant amount of good for the money that they spend. And so, we're really starting with these programs as one of these points of our focus areas. And so, all of these organizations do that.

03:41 EH: The first one, the Against malaria Foundation, funds bednet distributions that protect against the mosquitoes that transmit malaria in Africa. And they were one of the organizations we recommended very highly for a couple of years. And then, last year we suspended that recommendation, because... Due to our recommendation, they had taken in about \$10 million dollars, and they had struggled to find ways to get that money out the door. And we felt that they needed to show that they could agree to more distributions with more on-the-ground partners before we wanted to direct them with more additional funding.

04:16 EH: The major update of this past year is that they agreed to distributions that cost roughly \$8,000,000 in total with a new partner working in a new country in the Democratic Republic of the Congo and we felt that that was good enough evidence that they could get money out the door that we moved them back to our top charities list. The reason that we like the Against malaria foundation, relative to other bednet organizations, is that they are very focused on collecting high quality data about the distributions that they fund. And so, what that means is that when nets go to people, they take pictures to ensure that the nets actually get there and that the people receive them and put them up. And they go back later and survey the people who've received the nets to find out whether or not they're actually using them.

05:07 EH: And the data that we have from that is... Shows pretty well, that people are getting nets and using them consistently. And this is data that... We don't know of any other source that has data like this for monitoring whether people are using the nets they receive as part of a distribution.

05:24 EH: Now, AMF is still relatively new at implementing this type of program. They have two distributions that have reached a level of... That have gone far enough, that we have the type of data I'm talking about. So, we think that, overall, bednets are a great intervention. They have a lot of evidence. They're very cost-effective. AMF, when it has implemented its program, has done so successfully. But there's still relatively limited data, and they have a relatively short track record in the scheme of things.

05:54 EH: The next two organizations both focus on a program called de-worming. This is treating children, primarily, for intestinal infections that have very few and hard to measure short-term negative effects, but there's some evidence that treating children for these worms has significant effects on their long-term cognitive development. Most of the case for... Until this year, most of the case for deworming relied very heavily on a single randomized control trial that has been conducted in Kenya in the late '90s. One of the issues with this trial is that it found very strong effects. It found that students were significantly more likely to attend school, that when followed up with 10 years later, they had 25% higher earnings, but it happened at a time when there was significantly higher infections rates than there normally were in those areas. It was during an El Nino year that diseases are waterborne, and so after the flooding that was due to this weather pattern, the infection rates really spiked.

07:00 EH: And so a question that we had was how representative are the deworming organizations result of what happened in that study relative to in what would you expect to happen if there were significantly lower infection rates as is likely across most of the places that are being treated today? The big update this year, is that a new study came out that also followed up on children who were part of a randomized control trial of a deworming program, finding that test scores for the students who are treated were significantly higher than those for students in the control group.

07:35 EH: Now, this was in the roughly the same time and roughly the same part of Africa, so it's possible that this was also affected by the El Nino weather pattern. But the intensity rates were significantly lower and so this to us was in additional indication that the program has strong long-term results.

07:56 EH: The major update for the deworming charities was really a deworming as a program update as opposed to a major update with the two organizations. The two organizations are Deworm the World Initiative which doesn't implement programs directly but it vices and provides technical

assistance to governance on their program implementation. We think they do a great job. They have monitoring that shows that are people are not as great as we would always like to see, but it seems to show that the programs are implemented effectively. The main update with them is that they seemed to have limited capacity to absorb additional funding from donors. They have a target of taking in about \$1.3 million over the next couple of years and we hope that they can receive roughly three quarters of a million dollars this year, but not significantly more than that because when they hit that cap, they have fewer uses for additional funding.

08:52 EH: The other organization that does deworming is the Schistosomiasis Control Initiative. Schistosomiasis is one of the infections that affect children and we call them SCI, because it's a little bit easier to talk about in that way. They are an organization that has implemented deworming programs across Africa for a long time. The issues that we've had pretty frequent communication issues with them and what that has meant is that we found it harder to really nail down how they spend money, what specifically they're able to do with additional funding that they wouldn't have been able to do otherwise. We had pretty significant miscommunication with them a couple of years ago, which we wrote about in our blog where there were studies that we had been relying on to assess whether or not their programs were successful and it turned out that we had misread and they hadn't corrected us on those studies and it appeared... And the programs were not measuring their success in the way that we had thought they were.

09:55 EH: And so this is an organization where we've had some of these struggles over the years. At the same time, we have now been following them for more than five years. We've visited them in Africa. We visited them in London, where they're based, multiple times. We've been all over everything we can find and we've also found no evidence that anything concrete is going wrong. And so our view is that this is an organization where we have struggled to nail it down as well as we would like. We communicate with them significantly worse than with our other top charities, but overall we think they're an outstanding opportunity to accomplish a lot of good. And with this year with all of the funding that we expect them to receive from donors not influenced by GiveWell, we're targeting them to receive an additional \$1,000,000 over the next few months. We think they'll put that to good use.

10:46 EH: The final organization that's on our top list is GiveDirectly and they do something very straightforward. They try to get 90 cents for every dollar you give to very poor people living in Kenya and Uganda. They're the organization that we've had the easiest time communicating with and overall seemed to be the ones that are most monitoring and sharing information about their quality of their programs. The real update this year for GiveDirectly was nothing is specific to them. They continued to grow and move money out the door. They received about \$17,500,000 last year. They either spent or committed \$13,000,000 of that and that is a big increase from previous years.

11:30 EH: But instead it's about the other charities on our list seem like better options to us this year than they did in previous years. Last year, we had taken AMF off the list so that wasn't an option for donors and the evidence for deworming has, we think, significantly improved since last giving season. And so that has led us to have a lower target of what we'd ultimately like to move to GiveDirectly. Based on projecting how much money we expect to come in over the next few months, we're hoping to move them about \$1,000,000. That said, we think they could effectively take in up to another \$25 million and would successfully be able to move that out the door.

12:09 EH: That's kind of a brief update on our four top charities, what they do and the major

updates from this past year and I want to pause and turn it over to you to see what questions you have.

12:24 Speaker 2: For the Deworm the World Initiative you said that you think they only have \$1.3 million of more funding. Is that based on their projections or your projections?

12:34 EH: Yes. So I'm going to repeat questions for the sake of the recording. So the question is, is the \$1.3 million in additional funding they could take in, is that a GiveWell projection or a Deworm the World Initiative projection?

12:47 EH: This is their projection of the amount of money... I should say, unrestricted money that they could effectively use on what they told us is that there are many large funders who are interested in funding the scale up of deworming and so if countries are prepared to basically take money and do more deworming they think that there are large funders out there who will pay for that and therefore they're not coming to us and saying, "We could take more money and use it to scale up programs." The type of work that they're doing with this money is really two-fold. One is more open ended policy oriented research questions that could influence how deworming funds are spent down the line. The other is trying to go into some of the countries that where they don't yet have operations and don't yet have deworming programs to enable them to be in a position to implement those deworming programs in the future.

13:51 EH: Yup. Right there.

13:52 Speaker 3: So one thing I started doing in preparation for Christmas is directing some of my relatives and friends to GiveWell and then one of the complaints that I often get back and this might be answered on a blog like you said...

14:05 EH: Yeah, please.

14:05 Speaker 3: Anyway. Is that it seems like a lot of GiveWell top rated charities are addressing symptoms rather than the illness. So could you speak to that point?

14:16 EH: Yeah. So the question is that, there's this great thing you can do which is direct your family and friends to GiveWell and in the event that you do that and they say these charities seem to be addressing the symptoms of poverty. They're addressing things like health problems as opposed to the root cause of poverty which could be something else. How should you respond to that? First off, I think that it overstates, I think of these health problems not only as symptoms of poverties, but also of potential cause. I think that its quite difficult to... If you have a fever from malaria to go out and be productive in the way that one would need to be to, for economic development and economic growth. That said, and so I think like if you go looking for something that is the "Root cause of poverty". I think that is now something you are very likely to find.

15:18 EH: We have a blog post that we wrote many years ago called the "Root Cause of Poverty" or something like that that tries to directly address this question. There's a lot of debates about whether aid causes growth or whether aid harms growth and I think that the answer is that we don't really know how well it works. That said, the GiveWell top charities are trying to be a very specific thing and serve a certain type of donor. I think what we were trying to do when we started GiveWell was find the... Start with these criteria which are looking for strong evidence of effectiveness and getting

significant bang for your buck where you can really check and verify that our program has worked. And that proved quite difficult and I think that there are definitely opportunities to give that are outside the scope of these giving opportunities that we recommended and I think the challenge is, as a donor, knowing whether those other opportunities are successfully addressing the problems that they're intending to or are just another non-profit telling a really great story about what it's trying to do.

16:30 Speaker 4: I have something to say, too. That question and I have a question for you. So what I've come to believe or understand is that, when organizations do something that has an outcome of kids doing better in school or the parents don't have to stay home to take care of the sick kid or with GiveDirectly that they have found that the families that get the money end up with more assets and less depression, if these facts are true. So I see it as really giving people a chance to break the cycle of poverty that they wouldn't have otherwise. And the question I have for you is with AMF, is their data that proves that malarial rates go down in areas where they're distributing?

17:19 EH: Yes. That's a good question. So the question is, does AMF have data that show malaria rates go down? This is something that AMF started collecting in some of their first very large scale distributions. For a long time they worked on much smaller distributions of significantly smaller numbers of nets. They collected this data in their first distribution in Malawi and what they found when they looked very closely at the data they were receiving, is they weren't confident that the data coming back was an accurate reflection of malaria rates in the community and they believed, though I don't think they've fully gone through this all and written it up yet, that one of the causes is that there's two different ways that you can be diagnosed with malaria if you go into a health clinic.

18:03 EH: One way is the clinician looks at you and measures your temperature and says, "Well this appears to be malaria," and then they'll mark you down as malaria in the clinic records. And then the other way is that they use a rapid diagnostic test which gives a more objective measure of whether or not someone had malaria. And there were periods of time after their net distribution where many of the clinics didn't have sufficient supplies of these diagnostic tests, and therefore they didn't believe that the data they were getting back was an accurate reflection of what was happening to malaria case rates. In some of the country... In the country they're going into next, the Democratic Republic of the Congo, they're not intending to even ask the partner to collect this data because they don't think that what they get back would be very useful.

18:48 EH: Now it's worth having the context that AMF was absolutely unique in trying to collect this data. It would be... I don't know of other... There are places where people are gathering, collecting, and analyzing this type of data, but it's not common practice for a net distribution to take place and then there to be data available on malaria case rates. AMF was a standout in trying to collect something different, and I think they ran into this problem where the data not being high quality, and so I don't know to what extent they'll be doing it in the future, but it's not going to be happening with every single distribution.

19:25 Holden Karnofsky: I would add it's a general fact that we've seen that malaria case rate data just overall tends to be very noisy and unreliable. There's just a lot of reasons malaria cases can go up and down and there's also a lot of reasons that measurements can go up and down, so there's all these different ways of diagnosing it and some of them give different results. And sometimes it looks like the data's going up, it looks like the cases going up because the data collection's getting better. And even when you look globally at just trying to see how are we making progress on

malaria, it's a very hard thing to tell and there's a lot of noise. So I think it's good to try to collect that data when possible. I don't think it's the most meaningful data, and in some ways I think the data about whether the bed nets are still in use long after being distributed is maybe a little bit more meaningful.

20:10 Speaker 5: Has there been any recent news about whether the mosquitoes are developing insecticide resistance?

20:14 EH: Yeah, so the question is about is there any updates on insecticide resistance? So the way that the mosquito nets that AMF gives out work, is that they have insecticide on them, and they're not just intended to be a protective barrier against mosquitoes, but they're... If a mosquito lands on the net, the mosquito dies, and that's an important part of the effectiveness of this intervention. I think our best estimate is that if you had... I think, based on... They did studies where they compared insecticide-treated nets to totally untreated nets, and the untreated nets had about 50% of the effect of the treated nets, so it makes a big difference.

20:47 EH: Now this is a big concern, because insecticide resistance has been rising throughout Africa, and this is measured by looking at the mosquitoes themselves and seeing that mosquitoes, they find, are to some extent, resistant to the insecticide. The thing that still does not appear to have been found is a significant material impact on effectiveness in the field. Now to some extent, I think that may be due to lack of studies that would have found it if it did exist, but this is certainly something the malaria community is focused on. There are trials underway right now trying to assess the extent to which this is in effect. And so in our overall assessment of the cost-effectiveness of AMF's program, we're building in the possibility that some of what you get is only a protective barrier effect, not the full insecticide effect. Yep?

21:49 Speaker 6: This is more of a question on your mission and what you're hoping to achieve in the world, but do you have any concern about creating a more diverse group geographically and issue area-wise, as far as who you recommend to give to?

22:05 EH: Mm-hmm. Yeah, that's a good question. So the question's like are we trying to diversify our top charities, our recommended list? It's not something we're trying to do right now. We see ourselves as serving a particular type of donor, and what we want to do is serve that type of donor really well. And if we try to do too much, I think we'd end up with worse results across the board. When we started GiveWell, and you could find this is old plans that are on our website, our vision was the menu for charitable donors. It was the best climate change charity and the best disease charity and the best animal rights charity, et cetera, et cetera. And in early 2011, we saw that there was a lot of interest from donors to give to the organizations we most strongly recommended. And even though we had something like a mini-menu at the time, where we had some US-based organizations and some other organizations working on education or microfinance, there was very, very limited interest in those relative to what... The ones we were most strongly recommending. And so our hope is that the way in which more information becomes available about other groups is not via us and our staff, but by others who take on the mindset or this attempt to provide better analysis and cover other causes with it, all the way and back.

23:33 Speaker 7: This is a sort of related question. Are most of the organizations that you went through now, US or European-based, working in Africa or otherwise that is indicative of most of GiveWell's recommendations over the years?

23:48 EH: Yeah. So the question is whether the organizations we recommend are based in Europe or the US, as opposed to being headquartered locally. All the organizations we recommend are headquartered in the rich world. They all have... Not quite all, but you could go through them one by one, but many of them, most of them have significant local staff on the ground in the countries where they work. But the headquarters is in the US and Europe. A lot of that just comes down to two criteria that we need in order to do the work that we do. One of them is we need to have people that we can communicate with easily and get information from and so certainly there may be outstanding groups working on the ground but if they can't get information to us if we couldn't understand what they're doing we won't be able to recommend them. And those are certainly the type of... That's like one case for a false negative. Someone who's doing great work but not on our list.

24:50 EH: The other thing is groups need to be of big enough size that it is worth our time to research them. Because if you have a \$10,000 organization working... We visited some a few years ago in India. If you have a \$10,000 organization in Mumbai, this is not an organization that will be able to take in a significant amount of funding. And so for us to put a lot of time in to those organizations wouldn't end up being very worthwhile to our donors. We have a blog post we wrote called "Why we recommend so few charities" and I think it's important to note that we don't think those organizations we don't recommend are bad. We see our list as highlighting and recommending donors give to places that we are very confident in and we ultimately have much less to say about the groups that we haven't looked at.

25:39 HK: I think on that topic it can also be helpful to just keep in mind the full flow of the money and where it ends up. So we recommend AMF, but it's not as though the money goes to someone in the UK who then is kind of like determining everything about where it gets spent. It's more like the money goes to somebody in the UK who then negotiates with an African government and a bunch of other partners and then that money is going to fund a lot of people who are working for the government in the field as it progresses all the way to the field. The people getting involved in the distribution are maybe the district health workers or other, there's various mechanisms for community involvement. The situation often looked similar with deworming. So I think it's in some ways more accurate to think about it as GiveWell is evaluating how to get the money to the field in the best possible way and a lot of times what we're looking at is the next intermediary who someone is able to communicate well with us. That doesn't mean the money isn't eventually getting through to the locals.

26:39 Speaker 8: So at this meeting are you going to talk at all about the past top charities that aren't on the list this year or the ones you were seriously considering for the past half year, three quarters of a year that didn't make the list?

26:52 EH: Yeah. So the question is about other charities we've looked at in the past and are not recommending and others we looked at this year and don't recommend. There aren't many groups that we, the groups that are on our list this year comprise the charities that have been on our top list now for a few years. And the single group that was on our list most recently that we don't recommend now is Village Reach. We haven't checked back with them in a long time. We followed up with them basically post recommendation to the point where we felt like we understood what was happening with the program that we had recommended. The groups we looked at this year comprised the four standout organizations, which I'll talk about next, two other groups working on

programs where we didn't finish that intervention level assessment of their work and so we felt like we weren't able to take a view on how good those organizations would be.

27:46 EH: So these were an immunization campaign to prevent neonatal tetanus, maternal and neonatal tetanus. And then a program to treat lymphatic filariasis which is another neglected tropical disease. And with both of those we still are working on the intervention assessments. We did ask Nothing But Nets, which is another bed net charity, to apply and they declined to apply or participate in our process and we wrote a blog post about that at the time. Our best understanding of that was that they are very small they have four or five staff members and didn't feel like they had the capacity to participate. That's what they told us. But let me use that as a segue to talk about the four standouts and then we can go back for more questions on our four top charities, these four standouts or anything else about GiveWell that you'd like to ask about.

28:35 EH: So these are four organizations that we looked at this year. When we engaged with them we thought about them as potentially ending up at the top of our list and I think some of them realistically could be at the top of the list in future years, but we didn't feel like what we knew this year about them led us to want to put them there right now. Two of the organizations work on a program that aims to fortify salt with iodine. And this is a program that has very strong evidence that the children who are iodine deficient don't have the same cognitive development as people who have sufficient levels of iodine and therefore this program because you can fortify salt and reach many people with relatively low cost it becomes a very cost effective way to improve the... To help people worldwide. Even though this is a program that has been accepted, we've had a salt iodization program in the US for decades, there are still countries in the developing world where people are iodine deficient.

29:42 EH: The two organizations that work on this, neither of them is going around and literally trying to do the process of putting iodine into salt. Instead they partner with governments and the salt industry to try and pass legislation, advise on regulation and advise on quality assurance and monitoring to ensure that these programs go well. And so I'll go through each of them.

30:09 EH: One is called ICCIDD, that's an acronym for an even longer thing that you won't remember. And they're a very small network of mostly volunteer academics who... Many of whom have geographic focuses. And they don't put very much time into this, but they are, they believe that their advocacy to governments makes a difference to the likelihood that the governments implement salt iodization. And that the programs end up going well.

30:40 EH: We talked to many of their different, what they call, Regional Coordinators, this year and we weren't able to find a demonstrable impact that they had on salt iodization programs. Now, they're really small, and their annual budget before this year was around half a million dollars a year and so, it's also... It seems quite plausible that they may have had a very big affect, we just weren't able to find it. And so, we hoped that with the additional funds that they receive, the GiveWell influence donors, in particular, Good Ventures, the foundation with whom we share this office space, and Cari is the president of, has given ICCIDD \$350,000 this year. We would be happy to see them give more, that we don't have an explicit target for them. We think that may enable them to scale up in such a way that we can better understand the... In future years, what impact they've had. They could certainly be a top charity in the future.

31:40 EH: The other salt iodization program... Organization, or the program is part of a larger

organization. They're the Global Alliance for Improve Nutrition or GAIN. They run a variety of nutrition programs. Salt iodization is one of them. They were funded for the last 10 or so years, by a grant from the Gates foundation. The grant is up early next year and we don't yet know, whether or not it will be renewed. But we think it's possible that will not be renewed.

32:08 EH: They sent us... So, they are more... They are higher intensity than ICCIDD is. They were spending a few million dollars a year on this. But they still are working relatively far from the ground. And though they sent us data, we weren't able to get into it sufficiently to really know whether or not, they had an impact. And so they're another group that, we think given the strength of the evidence for this program and its cost effectiveness, we think they could end up being a top charity in future years. But we didn't, this year, have enough bandwidth to go through the, what we saw from them and really determine whether or not, we... We think they've had an impact.

32:49 EH: The last two organizations both have, either run or in the midst of running a randomized controlled trial of their own program. So one is called, Development Media International. They produce mass media that aims to encourage improved health behaviors. So, this could be something like, if your child has a fever, you should take them immediately to a clinic. If they're coughing it encourages excluding breast feeding. And the idea is, if you could actually change how people behave, because you can reach so many, so quickly, it would be an incredibly cost-effective way to improve health and save lives. Now, they're running a randomized controlled trial of their program in Burkina Faso, where they randomized some radio stations to broadcast, and some not. And they've measured significant changes on self-reported behavior change.

33:39 EH: So when they ask people, "How did you behave, when your child came down with diarrhea, they find positive self reports of people following the instructions that the spots laid out. The problem is knowing how to translate self-reported behavior changes into actual changes in behavior and we're hopeful that when we're excited about the final results that DMI will be publishing in the next couple of years, which will measure changes in child mortality, so it won't just rely on this self-reported behavior change. If... Even now, with us being skeptical about the evidence and trying to do our best projection of where we think they'll end up, our best guess is they are pretty cost-effective organization. But if they are... If they achieve the types of impacts that they expect to, they... They certainly have a good chance of being on our top list in future years.

34:34 EH: The final organization is, Living Goods. They're a group that manages and supports a network of what they call, "micro-entrepreneurs", who go household to household, selling health and other household goods. So, Living Goods will buy these goods, like a bednet, or a vitamin A supplement at... And we'll buy that, we'll sell them at wholesale to these individuals. And the individuals will go to door-to-door and sell them. Their randomized controlled trial completed earlier this year and they found a 27% reduction in childhood mortality among the group that was served by Living Goods relative to the control group and that's a very significant result.

35:17 EH: We're... The reason they are where they are on our list is that, when we do our calculation, our cost-effective calculation, they don't come out as... The cost per life saved seems is not as attractive as it is for our other top charities and so, that puts them down the list.

35:35 EH: The other big issue with them right now, is that the trial results were made available to us, but we're not able to talk publicly about the details of the study, because it's awaiting publication in an academic journal. Which is reasonable, but we're not able to talk about it all and so, that leads

us... That's another factor of Living Goods.

35:59 EH: Those are the rest of our list. One thing that was exciting for us this year, is that it is probably the most charities and the most research we've been able to do since GiveWell started, and a lot of that... We've been expanding our staff over the last couple of years, it can take a long time to translate new hires into new research production, but I think this was a year where that really happened. For a long time, this research was just Holden and me and a couple other staff members. And this year Holden was basically uninvolved in this work and my role was diminished significantly from where it had been in past years. And that was for us a really big organizational success to bring people on and have them contribute so meaningfully to our work this year. So, let me pause there again and open it back up to you all for any questions. Yeah, right in front.

36:55 Speaker 9: So, you've mentioned that the last charity didn't have a particularly good cost per life saved. What do the numbers look like?

37:05 EH: So, our best estimate is...

37:07 HK: Can you repeat the question?

37:09 EH: Oh, yeah, probably, yeah. So, the question is about Living Goods... I'm trying to adjust for how far people are from the mic, but it's a good reminder. The question is, what our best estimate was for Living Goods cost per life saved. Our best estimate is around \$11,000 per life saved. That's about three times as much as our estimate for the Against Malaria Foundation. That said, I think it is really worth being cautious about how you interpret these figures. It is... You can... We have the Excel files online, you can poke around with them, you can see what goes into them, and they rely on a lot of hard to estimate inputs and also philosophical value judgments. We use this as one input into our process, but certainly not the only one, and I think it's worth... We look at that as some indication of when we put every charity through a... We're doing like more or less similar things through this same framework. This is one factor that we can use in deciding between them.

38:10 EH: It did end up playing a pretty big role in our ranking or our recommendations for our top charities, because GiveDirectly is an organization that we, collectively on staff, feel really great about. They are excellent at monitoring and sharing information, both things that are good and bad, and they seem to have scaled up over the last few years to reach many people. On the other hand, SCI, which does deworming, is an organization that we've really struggled to understand. It doesn't have the same type of monitoring data. But our best estimate of the relative cost-effectiveness is that SCI is, or deworming, is about five to 10 times as cost-effective as cash transfers. And so, with that multiple in mind, I think it led a lot of people to feel that basically that difference in cost-effectiveness made up for the differences in organizational views.

39:04 HK: So, back on Living Goods... Their cost per life saved is a bit worse than our top charities, but maybe competitive with GiveDirectly. But it would have had to be a lot better to get them on the list, because they have much less of a track record, they have one randomized controlled trial that was their own study. It's one case. There's not as much that we know about them, so I think for the stage they're at they would have had to be a lot better. I think DMI has a pretty competitive cost per life saved, depending on how you slice it, but also less of track record, only one study, problems with the study, et cetera.

39:41 EH: Let's go here, you next.

39:44 Speaker 10: I imagine that when you do try to measure cost effectiveness, you do have to deal with the problem of how you discount the future. How does X in a year compare with Y in 10 years? How do you deal with that? What's your discount rate and how do you evaluate it?

40:00 EH: Yes, so the question is: How do we deal with discounting future effects for the impacts that charities have? One is a very mechanical. The way in which we model cost-effectiveness for three of our four top charities. GiveDirectly is giving people cash that enables them to have an additional stream of income over future years due to investments that we think people tend to make when they receive large inflows of cash, and we discount those cash flows back to have the present values of that money. We do the same with deworming. The main effect that we look at for deworming programs is the future earnings that people have, so we discount those. Those are done in a similar way.

40:52 EH: For the health interventions, we're not applying an explicit discount rate, we're saying "In our view, when this program happens, it will prevent this person from, or this number of people, per dollar, from dying from a disease in the near future", and we think those are... To me, those are the most comparable items that I can hold in my head. We don't try to put everything in a single formula, where we say, "This is the value of this person's life" over some long period of time and try to compare it directly to the cash.

41:26 HK: I would recommend if anyone's really interested in our cost-effectiveness, you can download our spreadsheet, It's kind of a mess, but if you play with it, [A] you can learn what judgment calls we're making and how you would make them, and we've tried to lay that out really clearly. And [B] you can just get a sense for how rough these estimates are and how you feel about them and how much you want to weigh them. One of the things we have in that spreadsheet is the discount rate. So, you can put in 0 or 10% and see what happens. We generally use positive discount rates, 3%, 5%, and I think there's a bunch of reasons that one should use some a discount rate, one reason is that we think that when you help people today, those people can then help other people, and those people can help other people so it's a little weird to just count. That's one of several reasons that's it's a little weird to just count a benefit 10 years from now the same weight as a benefit today, but we don't use huge discount rates either.

42:23 S8: It's seems like in many of its funding priorities, GiveWell stays away from projects that seem to have high expected value, but maybe are more risky bets, so this might look like research tracks to sterilize the vast majority of mosquitoes that are causing malaria or some of the far future causes represented in the Effective Altruist movement. So what keeps GiveWell staying away from these sorts of causes?

42:46 EH: Yeah. So the question is GiveWell is focused on I guess, shorter term, more direct impact causes and not more speculative issues. So that is a good segue into our next session, which is the Open Philanthropy Project. That's the part of our work where we're focusing on potentially high impact philanthropic opportunities that don't have a strong track record of success, where you can't see what had worked in the past and then directly try to make that happen again in the near term.

43:18 HK: And I think as we do that session, it will partly answer that question because the

experience I've had now that we have the Open Philanthropy Project rolling pretty well and we have sessions on and everything, it's just like... It's way... It's harder to explain why we believe what we believe. It's harder to make the case, it's harder to get people onboard, a lot of people come in and they hear what we're saying and they say, "That's sounds wacky," And they would take me maybe days of taking one of these ideas and trying to understand it and understand how you feel about it.

43:49 EH: So, I think there was good reason and there is good reason when serving an audience of individual donors, people who have a few hours a year to think about what they're going to do, I think it makes a lot of sense to focus on these top charities which well there's certainly a lot of information, and the information is pretty concrete I think it's possible to see how we're reasoning, how we picked these charities, why these charities. What you can expect to accomplish by giving to these charities. I think that becomes a much bigger lift with Open Philanthropy Project and that becomes a better fit for people who want to spend a tone of time thinking about where to give instead of some time.

44:23 Speaker 11: Yep. I think someone answered this already, but I think it's good obviously that you guys have the grand vision and you started with kind of a pragmatic approach with what you can do today, but you guys are part of something larger which is creating quantifiable metrics to evaluate outcomes/outputs across the board. So my question is: As founders and as your journey as GiveWell, have you guys found this to be more of a competitive space or more of a collaborative space in terms of the other players and what they're evaluating what you are evaluating, is it open or is it combative?

44:57 EH: So yeah, so the question is about, I guess, how we GiveWell fits in with the other players in these space and how we interact with them?

45:06 S1: Yeah.

45:07 EH: I think the short answer is it's not clear what space we're in and who else is doing this type of work.

45:13 S1: Yeah.

45:13 EH: I think about what we do, that what we do is very unique and we haven't really found other places that are, or doing something similar enough that they're like clearly a peer group. When we started the peer group was the charity evaluation groups, like the Charity Navigators of the world. We talked a lot to them early on but their... And I don't mean to pick on them but they're the most well known and they're certainly other groups in that basic framework. I think that what we're trying to do is sufficiently different than what Charity Navigator or Guide Star is trying to do, that we, given where we are now we have less to collaborate with them on, that can really be helpful. So Charity Navigator and Guide Star are trying to be a place where you can go to a website, type in a charity's name and get some information about it and you really have to have that charity in mind, and you're trying to check very basic things like, how much does the CEO make or are they spending it all on fundraising or administration.

46:13 EH: And we're serving a donor who says, "I want to accomplish good with my charitable donations, I don't know where to give," And they're looking to someone else who's put in that work to figure out which organizations really get significant bang for their buck that they can rely on it.

46:32 HK: When I think about the Open Philanthropy Project, the biggest competitors or something, I usually think about foundations just 'cause there are other groups that are trying to give away money really well and we're often sort of saying, "Well we would do this," But then it turns out someone else is already doing it and blah, blah, blah, although I think they're very collaborative and I don't perceive them as competitive at all, and I think they're trying to help us. I really cannot think of another group that I think is doing a similar service to GiveWell in terms of finding top charities, doing the deep digging on charities like this.

46:59 EH: One more question and then we'll break. Yeah, go for it.

[background conversation]

47:05 Speaker 12: I was just curious along the same lines as you've been doing this now for five years, have you find it either the research process is getting faster and more efficient and/or are there are field-wide norms that are growing up around the kind of information you're looking for that maybe influencing other organizations even if they are not ending up in the target list. Do you see evidence that you are actually incenting shift in behavior around accountability and transparency?

47:30 EH: Yes. So the question is: How has our experience of doing this work changed over the last, the years that we've been doing this? I think there's a few ways in which it's changed concretely. One is that the most obvious change is we are better known and the incentive to organizations to participate in our process is higher and that has made our job much easier. When we started GiveWell, we would call up charities and say, "Would you talk to us?" And they would just, would have no interest in spending time with us. Obviously because we were going to take a lot of their time and the best they could hope for at the end of the day was \$5000. That wasn't a great deal for the charities.

48:06 EH: We also were much more prone to writing crazy things on our website and blog that bothered some of them. I do think that there is some movement on the monitoring, evaluation, transparency front, though it's really unclear whether or not that is due to GiveWell or coincident with GiveWell. The first organization we ever saw that had run a randomized control trial of its own program was GiveDirectly, that came out last year, there's two more that came out this year. I think what we want to do is to be there to assess these trials when they come out, assess the organizations. So if you do go through that process and you succeed you will have the reward of significant funding from donors. And that's... That kind of mechanism didn't exist before and we hope that to some extent we provide it.

48:56 EH: Alright, so let's pause here. We'll take a five or 10-minute break for people to use the restroom or grab some more pizza and then we'll come back for the more, crazier session.

[chuckle]