

GiveWell San Francisco Research Event June 6, 2017 - 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.

00:02 Catherine Hollander: Hi everyone. I think we're going to get started now. Thank you for coming. I'm Catherine Hollander or Cat, nickname there, a Research Analyst within Outreach Focus at GiveWell. And this is Elie Hassenfeld, GiveWell's Co-Founder. And we're going to be co-presenting on GiveWell's work tonight. Just for a brief overview of the event, Elie and I are going to talk for about 50 minutes leaving a lot of time for questions in that period, focusing on the work that GiveWell is doing to find top charities. Then we're going to take a quick break, you can get some more food, and then we'll come back after about a 10 minute break and Holden Karnofsky will talk about the work of the Open Philanthropy Project then. So, that's kind of the general overview of tonight's event.

00:51 CH: Specific to GiveWell, we're going to try to focus the talk as much as we can on new work that GiveWell is doing. But I'm going to start with just a quick overview of what GiveWell is. We really hope that if you have any questions about GiveWell just in terms of what we do, who we are, that you'll ask them during the Q&A. We don't want you to feel like a question is too basic, or you can't ask about just like GiveWell in general. So, that will be something we encourage during the Q&A. So like I mentioned, I'm going to talk about kind of a brief overview of GiveWell and some organizational updates. And then, Elie is going to talk about GiveWell's Incubation Grants program since he's leading that work and research. So, one other just housekeeping note as we always do, we are recording this event so that we can publish it on our website. If you say anything or ask a question that you would like removed from that recording, that is totally fine. Please just come up and let me know or send an e-mail to info, I-N-F-O @givewell.org after the fact, and we're happy to take that out before we publish it on our website.

02:02 CH: Great. So, I'm just going to start with the sort of quick overview of what GiveWell is and what we produce. GiveWell is a non-profit dedicated to finding outstanding giving opportunities and publishing the full details of our analysis to help donors decide where to give. So, we publish an annual shortlist of our top recommended charities that's available on our website for anyone to use as well as the details of why we recommend the organizations that we recommend. We currently have seven top charities, and these top charities were announced as we typically do at the end of last year. So during giving season when most people in the United States make their donations. So, just quickly, the seven top charities that we recommend. Two of them work on interventions to prevent malaria. One of them is the Against Malaria Foundation which supports the distribution of insecticide-treated nets to prevent malaria. And the other is the Malaria Consortium which works on a variety of programs, and we recommend it just for its seasonal malaria chemoprevention program which is the distribution of preventive antimalarial drugs to children that are between three and 59 months of age.

03:17 CH: We also recommend four charities that focus on deworming programs. So programs to treat intestinal parasites. Those are Deworm the World Initiative and the Schistosomiasis Control Initiative. And then, two other larger organizations that work on a variety of programs including deworming, and we recommend them for their deworming programs, and those are Sightsavers and The END Fund. And then finally, we have one other top charity which is called GiveDirectly and they distribute direct cash transfers to very poor individuals in East Africa. So, those are the current group of top charities that we have. We're following all of their work and within I toured, updating our recommendations at the end of the year. So, we don't have any kind of major mid-year updates on the list of top charities. The primary thing that's changed on the top charities front at this point in the year is our recommended allocation.

04:12 CH: So, GiveWell publishes, for donors who feel highly aligned with us and want to follow our recommendation for where among the seven top charities that we recommend, we think that the donations can do the most good. We, at the end of last year, had recommended giving 75% to the Against Malaria Foundation of your donation, and 25% to the Schistosomiasis Control Initiative. We are now recommending that you give 100% to the Against Malaria Foundation, and that's a change that we made in April. And we are next planning to update that in November but if we have new information or our assessment of where the highest priority funding needs are, we will update that sooner.

04:55 CH: And the other piece of information related to that is that we made a grant of the donations that we received that were given to GiveWell for granting at our discretion. So, if you come to GiveWell's website to make a donation, you have an option to give the money to GiveWell to grant it where we believe it will do the most good, where we think the highest priority funding needs are. And we granted the funding that we received from giving season to the Spring add to 4.4 million to the Against Malaria Foundation and half a million dollars roughly to the Deworm The World Initiative. So, those are sort of the two updates that we've had on the top charities front. And otherwise, we're planning

to continue to follow and assess the organizations that we recommend. And as Elie will probably mention, look at a couple new ones for through the GiveWell Incubation Grants program as potential top charities. And if Elie doesn't mention it, we can talk about it in the Q&A as well.

05:48 CH: And then finally, just a couple quick organizational updates for GiveWell. We're now seeing outreach as more of a limiting factor than research for GiveWell. So, GiveWell is about 10 years old, or is 10 years old this year, and we have not focused a lot on outreach in our history. And we, last year, estimate that we left over \$100 million in high priority funding gaps unfilled, and that's excluding funding gaps that give directly. And so, we think that there is a strong reason to prioritize putting more into outreach. So, getting more people aware of the charities that we recommend, and hopefully moving more money to those since we now see a very large unfilled funding gap at the charities that we recommend.

06:32 CH: We don't yet know what outreach strategy we're going to adopt to do this but happy to kind of talk more about that in the Q&A. So that will be something that we're thinking about in 2017. And then the other organizational update is that GiveWell and the Open Philanthropy Project are now separate organizations as of last week. We're planning to announce the details of that this week, and we'd love to be able to announce it on our blog. So, we ask if you can please not tweet this, we definitely appreciate it. And that's kind of the other organizational update so far this year. And now, I'm going to turn it over to Elie to talk about the GiveWell Incubation Grants and other research work that we're doing.

07:15 Elie Hassenfeld: Great. Thanks, Cat. Hey, everybody. So, I would love to tell you about GiveWell Incubation Grants. And GiveWell Incubation Grants are programmed to try to support the development of future GiveWell top charities. And there's basically three types of grants that the GiveWell Incubation Grants program is looking to make. The first is organizations that could potentially be GiveWell top charities but they need funding to get off the ground. And so, the start-up funding is one type of grant that we'll make. The second is monitoring and evaluation support. So, you could imagine an organization that's running a program that could be exceptionally effective and cost effective but it doesn't have the funding or the know-how to implement a high quality monitoring and evaluation of its work. So, we're looking for opportunities to provide funding to organizations so that they can do that. And then finally, perhaps at the earliest stage, is grants to evaluate the effectiveness of programs.

08:25 EH: So, to step back for a moment, the way that GiveWell research process works is there's really two prongs, there's assessing the evidence for a particular program. So for example, how well does deworming work in general, how effective is it, how cost effective is it, and that's independent of the particular implementer that's delivering the deworming program. And then separately, there's the assessment of the implementing groups that one gives money to and they go and implement the deworming program. As part of our process, we tend to start with the program-oriented research because most charities don't have the resources to run a high quality study to determine whether the program they're implementing has the impact that they're hoping to see. What we have often relied on are randomized control trials to determine whether or not a program is effective. And so, this is the third part of the GiveWell Incubation Grants program, the funding to really fund the initial research to determine whether or not a program has the sort of effect that we're looking for.

09:30 EH: And the reason that we've decided to do this is, GiveWell is looking for organizations that have a track record of delivering a program effectively, and cost effectively. Organizations that are, whose programs are effective, that rely on existing research that demonstrate that those programs really work. And then, organizations that have that monitoring and evaluation in place so that we can see whether or not the programs are truly working. And so, in all of those cases, some other funder had to come in at some point to enable an organization to be eligible for a GiveWell recommendation. And with that in mind, we would like to be there to say, "If an organization seems like it's potentially promising as a top charity, we would like to provide down the line with additional support. We want to provide the funding to potentially enable it to get there." The Good Ventures is a foundation that you may have heard of, they're a group that we work very closely with, they've given significantly to our top charities in the past. The benefactors are Dustin Moskowitz, who's one of the co-founders of Facebook and his wife Cari Tuna, and they funded all of the GiveWell Incubation Grants to date. So, we've made recommendations to them and they've provided the funding.

10:47 EH: So, I want to just go through a few examples of the sort of work we're doing in the Incubation Grants program. So, one type of work that I mentioned was this improved monitoring and evaluation. And so, an example where we're doing that is surgical programs or an area that I think are very intuitively appealing to a lot of donors. There are claims that maybe some of you have seen that I was certainly very attracted to in GiveWell's early days that are along the lines of \$20 enables you to perform a surgery that corrects blindness for someone who has severe visual

impairment, and that is an impact. \$20 for a blindness reversed that would, I think if it were true, would be extraordinary compared to the other charities that we're recommending. When we looked at organizations performing the program, and often the program is cataract surgery which is relatively simple, relatively straightforward to conduct, they very rarely have the sort of monitoring and evaluation that we were looking for. I shouldn't say very rarely. We never found an organization that had the monitoring and evaluation that we were looking for.

12:02 EH: And to be clear, that's not to say that we knew they were not effective but we didn't see the evidence that they were. So, the sort of questions that we had before the surgery, what's the level of visual acuity? How much vision do people have? Because there's the minimum standard at which an organization might perform a surgery is impaired vision, but I think nothing that we would normally consider close to blindness. But then all the way at the other end of the spectrum is severe visual impairment and organizations didn't have useful information that helped us answer that question. Similarly, we weren't able to find consistent information for how much vision improved after the surgeries. So, how often was the surgery successful? How much visual acuity improvement did the organizations programs create? And then finally, there's just this question of if you're a donor and you give money to an organization that is aiming to increase the number of surgeries that are performed, does your money truly lead to more surgeries occurring that wouldn't otherwise have occurred? Which is really the question that I think donors should care more about. You can give money and the question is, what's the counterfactual? What would have happened if you didn't give that money?

13:20 EH: And cataract organizations run a wide variety of programs. They will do demand generation, where they try to get the word out that they're performing surgeries to encourage people to come in. On the other end of the spectrum, they'll partner with hospitals to provide general support to hopefully enable the hospital to perform more of the surgeries that it could. And so, we have partnered with IDinsight. IDinsight is a group that focuses on high quality evaluations to inform funders and other policy makers. And we're conducting a project where we're trying to identify a small set of organizations focused on cataract surgery, where we think that the combination of putting in place a monitoring and evaluation system to measure pre and post surgical conditions and outcomes, and then also trying to emulate some sort of impact evaluation to determine whether additional funding leads to additional surgeries being performed, would... And so, we're going to try to run that project and then we'll be in a position, potentially, to recommend cataract surgery organization. So, that project has just gotten started. We've written a little bit about it on our website but that's an example of the sort of monitoring and evaluation project that we might run as part of the GiveWell Incubation Grants program.

14:45 EH: Another example is trying to gather more research that could inform our recommendations. So, as Catherine mentioned, one of the programs that many of our organizations implement is mass drug administration to treat parasitic infections that affect hundreds of millions of children in the developing world. And this program is actually fairly controversial because there are academics on one side that say, "The evidence that this program is effective is very minimal and this program is not worth a significant investment." And then on the other side, you have another group of academics who say, "This program is extremely impactful and is one of the best bets." And we find ourselves somewhere in between... I guess, our bottomline is not in between. We certainly think deworming is one of the best opportunities for giving but we see significant gaps in the evidence base, and we are highly uncertain about whether or not deworming programs have significant impacts. We recommend it because the potential impacts are so large that we think it offers really good value to a donor who's operating in an expected value framework.

16:03 EH: But the question obviously arises, what could we do to get better information about how effective deworming programs are? Unfortunately, it's very difficult today to run a new randomized control trial that evaluates whether deworming programs have the sorts of effects that we're looking for. And just to be clear, to sort of layout what we know and what we don't know, we know with very high confidence that when you treat children in areas that have high levels of worm infections, the pills effectively kill the worms that are infecting the children. The question is, how much of that matters to one's long term quality of life? And there's a single randomized control trial that followed children over a very long term, 15 years later, and it found that the treatment group had significantly higher consumption. They were able to spend significantly, more presumably because they were able to earn much more than the children in the control group. Now, this is the only study of the deworming program that has followed children long term. And so, is in many ways a unique piece of evidence that points to a potentially very large effect.

17:16 EH: On the other hand, there has been many other attempts to follow children over a relatively short term and these studies have found only moderate effects on indicators like weight, school attendance. There's been no evidence of effects on improved performance in school. And so overall, I would say the evidence for deworming is somewhat murky. Okay. So, happy to answer questions about that if that was not fully clear. But suffice to say our plan now, what

we would like to do is to improve the information we have about these programs. Now, we don't think we can fund a new trial starting today for two reasons. The first is just that it's going to take a really long time for those results to come in. If you tend to deworm kids in elementary and middle school, so ages five to 15, by the time they enter the workforce, and then you get some data on their earnings. In many ways, you're talking about results that will only be available 10 to 15 years from now. And I think there's a reasonable argument that as the world continues to get richer, and I believe more funding flows to more effective opportunities, by the time the research comes in, it's not even clear that deworming is going to be a problem that will require more funding anymore. I certainly hope that it doesn't.

18:35 EH: The second problem is, that one of the reasons I think these effects might have been detectable is that the infection rates among children that were part of the study that I've described as having these long term effects and earnings were extremely high and much higher than they are in most locations today. And for that reason, running a study that could get the statistical power to potentially detect these effects would almost certainly be so costly to not be worthwhile.

19:07 EH: So, the thing that we are trying to investigate, the possibility we're trying to investigate, is looking back at some of these initial trials of deworming programs that only followed children over the very short term, and trying to follow up with them 20 years later. And so, there was a study... There were all sorts of deworming trials that were conducted in the 1990s. And if we found a study where we were able to follow up with those students today, we might be able to gather additional information about how effective this program is. And so, that's the type of study that we're investigating as part of our Incubation Grants program because it could significantly affect our ultimate conclusion about how effective we should consider deworming programs to be. This actually is something we're thinking about doing in other cases as well. We at GiveWell care most about the long term impacts of programs. And unfortunately, there are very few studies that follow participants over a long term. Many studies find effects but they're over the course of one or two years. And I think it's worth questioning whether effects on test scores from an education program lead to better life outcomes in adulthood. Or effects on weight or height from a nutritional program in childhood lead to improved life outcomes later on for someone's life.

20:32 EH: And so, we're considering a project that tries to more systematically look for trials that we could potentially follow up on long term and run a long term survey of the treatment, and the control groups to see whether the programs that were only followed short term actually had long term effects. Just want to quickly run through some of the other things that we're doing as part of this program just in case it raises some ideas for questions. So, there are a few other groups that were as part of the GiveWell Incubation Grants program, one is called Zusha! . They put stickers in the mini vans that many people take in Kenya that encourage passengers to yell at the driver if the driver is driving unsafely. Car crashes are one of the leading causes of death and disability in poorer parts of the world, and this program was evaluated in multiple randomized control trials where they found the significant reductions in accident rates.

21:34 EH: So, we've provided them some funding and are watching them closely. We've also are looking at two groups that are trying to increase utilization of routine immunization services, and both of these groups were set up with the explicit intention of trying to be a GiveWell top charity. So to get that start up funding and then make it to the sort of end and receive the recommendation. One is running a conditional cash transfer program to encourage caregivers to bring their children for the routine immunization schedule in Nigeria. They are called New Incentives. And the other is called Charity Science: Health. They're trying to run an SMS reminder program. So they will collect the numbers from families when their child is born and then send them a text message reminder over time to convince them to come in. These are both fairly early stage organizations but ones that we're looking at.

22:27 EH: So, there's a lot else that we're doing but I'll just say, the overarching goal right now of the Incubation Grants work, but also the ongoing top charities work, we're evaluating organizations that have a track record, and then the goal that Catherine mentioned of trying to raise more funding for our top charities is that we really want to create the incentive for people to start and run and operate demonstrably cost-effective programs. And we feel like we're really uniquely positioned to evaluate and provide funding really across the pipeline where we can say to someone, "If you want to start something, we can give you start up funding. If you need help with the monitoring, we think it's promising, we can do that. We can fund the research. And then, we also are able to influence a significant amount of funding when you reach scale so that you can also see your end game."

23:18 EH: And so, where we've provided or we've recommended small amounts of funding through the Incubation Grants program last year about \$10 million, our top charities collectively received more than \$90 million. And so, we really hope that the work that we're doing across the spectrum creates this funding mechanism and incentive that gets people to try and create the best programs that they can to improve the lives of people in the developing world. That sort

of ends the talk and now we want to turn it over to you to ask any questions that you have about things we've said, about things we didn't say. And like Catherine said, we know our website can be somewhat dense and hard to parse. And so, please use this opportunity to ask basic questions if you have them and don't worry about whether you think it's too simple. Yeah.

24:08 S3: So, this comes to mind is the idea that you're zeroing out your giving in that area because of the lack of reliable data or metrics to continue. How does that impact the organizations that are doing this work and their ability to carry forward, and the people who would otherwise be receiving some form of treatment?

24:36 CH: So, what happens basically if we stop recommending funding in a particular space? And I'm also going to, and Elie will, repeat questions for the recording so that they're audible. Yeah, I think Elie actually might be the best person to speak to what happens when an organization no longer is a top charity question since we've had a couple of that in the past.

24:55 EH: Yeah. So, our goal is that every year, our top charity list represents the best organizations that we know of that we're recommending funding to. And so, to keep that possibility in mind that we could change the list, what we tell our top charities is that when they receive funding from us, they really need to treat it as multi-year funding rather than funding that they should spend down in one year. And that's our goal. Now in practice, I don't think they've always done that. And I think that they often anticipate that we'll renew the funding at the same level. And so, this remains I'd say, a communications challenge with them that I don't think we've fully resolved.

25:38 CH: Yeah?

25:39 S4: I'm a donor to GiveWell. Your approach really appeals to me. However, a lot of people find it, not just not compelling, but offensive. I'm curious in the outreach, how you're going to address that. Are you going to be looking for more cerebral people or make a less cerebral approach?

[laughter]

26:02 CH: Great. So the question is, how are we going to address the fact that GiveWell's approach to giving isn't appealing and might even be offensive to some in sort of a broader outreach push. I'd be curious if you have any like specifically the offensiveness would be sort of the quantification side of it.

26:20 S4: Well, I saw a headline, "Giving to Little League Doesn't Make you a Monster," criticizing your approach.

26:31 CH: Got it. Yep. So, I think with outreach, do we think something like, "Giving to Little League makes you a monster?" or "How do we talk to people who might give to things that are different from what GiveWell recommends?" When you think just in... I guess, I'll answer this first from a personal perspective of how I talk about GiveWell to people, including friends and family who give to really different things than I do sometimes. And then also, what does this mean for GiveWell's outreach strategy going forward? And just first on the sort of personal aspect, I definitely agree that giving is personal. People give for all kinds of reasons. Sometimes people give to have the biggest impact they can, sometimes people give for social reasons, or to connect in their communities. And it's not that our goal is to make sure that no one ever gives in any other kind of way of recognizing. I think I do and other staff have talked about recognizing different ways that people give. And the type of giving that GiveWell staff get excited about and I get excited about is giving such that I can have the greatest impact with my dollar. And so, I tend to just focus on the fact that that question makes me really excited when I think about where my money is going.

27:45 CH: And then in terms of what that means for GiveWell's outreach, I think that there are... I mentioned, we're kind of early in thinking about our strategy and there are many ways that it could go. I think there could be a world where we just try to find everyone who's already thinking about giving in the way that we are, 'cause we often hear from donors like, "Hey, I was thinking about giving to maximize my donation and I was Googling, and that's how I found GiveWell, because I was asking the same questions that you were." And I think there are probably more people out there who are thinking in this way but maybe don't know about GiveWell yet. So, could we just kind of try to find more people like our current group of donors by reaching out to that group?

28:22 CH: I think another approach would be to sort of try to encourage people to think about this question that we spend all of our time thinking about which is, if you're giving, what's compelling you to give? And if it is an impulse to improve lives the most or have the biggest impact possible, that there is a resource that exists to help you answer that

question. And those would be two different types of outreach approaches, I think. And I don't think we know yet for sure which strategy we will take. But my impression is there are many people who are kind of already thinking in this way. That might make sense to reach out to starting out with since they're sort of approaching giving in a way that we are, as well. I don't know if Elie has anything else.

29:02 EH: Yeah. I think there's a... We don't have the ambition for GiveWell to be a mass market product where we influence how 50% of US donors, or donors around the world where our donors come from, give. And instead, I think if we can... People pull off and say this sort of thing, but if we influenced 1% of US charitable giving, that would amount to a few billion dollars. And so, I think the most likely approach we will take with outreach will likely still be on the cerebral side of the spectrum where we're really trying to make sure that all the people who would potentially be interested in us know about us and understand our value added.

29:51 S4: Thank you.

29:54 S5: It seems like a bunch of groups were gung-ho about deworming in prior years, right? And so, now you're saying, "Oh, well. We believe it's maybe... There's no proof that it's as good as people necessarily think it is." So that's good, increasing the accuracy of the predictions and stuff like that. But it raises the question like, "Okay. The implication is, some stuff was wrong for the past 10 years." It's like, "Oh. What else was wrong, right?" So, I just was wondering if you could characterize, as of today, how do you view your confidence in the accuracy of the estimates?

30:37 EH: Got it. Can I...

30:38 CH: Yeah.

30:38 S5: And is that broadly improving over time or...

30:42 EH: Yeah. So, there's three different way... So the question is just about the deworming and how it seems like that our view and other's view of the program has evolved over time. What does that mean for all the other programs and our views of them today? So, I think there's three important points to make. I actually think that GiveWell's perspective on deworming holds up very well. If you look back at our take from 2007 and our view has not updated substantially. I know that we did not communicate about deworming as well as we should have. And there was a time when we talked about our criteria as proven cost-effective and scalable. And several years ago, we decided to get rid of the "proven" word and now we say "evidence-backed" to really try to get away from the notion that GiveWell's top charities are certainly accomplishing good because there are a lot of questions.

31:40 EH: But we could, if you want, you can sort of see this on the GiveWell website or I could walk you through it. But I actually think that we have not changed our minds very much about deworming and you can see the uncertainty baked into that, historically. Five years ago, I wrote a blog post called, "Why I Didn't Give To SCI This Year?" SCI being one of our top charities that recommends deworming programs. So, the second part of that is how would I characterize our view today? And I would say it's a view that it's high uncertainty. We are in a situation where there is one trial that found very large effects on deworming. We have done everything we can to poke at that trial, to look at some related trials, and it holds up extremely well. David Roodman is an economist who's on the staff of Open Philanthropy. Last fall, he did a deep dive into the evidence for deworming.

32:41 EH: This is obviously considered the source where this is coming from but I think he may be the best in the world at trying to evaluate and synthesize evidence. This is not an activity that's very interesting to many people but David is great at it. And usually, when David looks at a study, he destroys it and finds the fatal flaw. And with deworming, it held up very nicely. And so, it's a challenging situation when there's a very, very strong piece of causal evidence but it doesn't match up very well with the limited short term evidence for deworming. And then finally, in that picture, we look at this program and we say, "Well, it costs 50 cents per child per year." And even if you discount the possible effect. Even if you say there's only a... Which is what I have in sort of my model. You say there's only a 20% chance that this effect on deworming is really there, that it's having this sort of large effect on earnings, it still ends up being more cost effective than the other programs we're recommending.

33:46 EH: And so, it's that expected value calculation that's driving the ultimate result. Now, I think it is worth you're taking that with a huge grain of salt. That 20% figure is largely a best guess, that is not particularly data driven, and I think it's... I wouldn't stand here and say that it is a mistake if you give to something other than deworming because here's our calculation that shows that it's effective. But our view is high likelihood a very limited effect but some

likelihood of extraordinary effect, and that puts it at the top of the list. I think in some ways, to the final piece of your question about what this means about everything else. I think I could tell other stories about work that we have done that should lead you to question the reliability of GiveWell's research. I don't think deworming is that case in particular, but we work in an area where often the data is very unclear, it's very unreliable. Overall, the world of, I would say, global health and development broadly but certainly charity in global health and development it's... You're just trying to work with extremely limited pieces of information, and that means that sometimes views can shift around significantly. So, just one concrete example of the type of limited information we have.

35:09 EH: The WHO and the Institute for Health Metrics and Evaluation at the University of Washington have significantly divergent estimates for the number of people who die from malaria every year. And they differ today by a factor of 1.6 or 1.7. They differ by more historically. That's a significant difference. If you think about, "Well, that's roughly a factor of two difference between these two organizations focused on evaluating how many people just..." Some very basic quantity, how many people die from a particular disease? And it has a major impact on how cost effective you should expect a malaria program to be. One of the goals we have with our work, with the Incubation Grants program is to generate more of this data to get better information, but I certainly think the state of the world today is the possibility of significant shifts as we get more and better data. Someone all the way...

36:06 CH: Right there. [chuckle]

36:08 S6: Are you having trouble with finding people to run these 15 year RCTs? And if so, are you communicating with organizations like MITs, Poverty Action Lab?

36:19 CH: Do you want to talk about that?

36:20 EH: Yeah. So, the question is about how we're finding people to run the long term RCTs. Yeah. We're not doing this on our own. This is talking to those people. So, talking to Rachel Glennerster at J-PAL, the Poverty Action Lab. Here in Berkeley, there's the Center for Effective Global Action CEGA and Ed Miguel there. I think they are the people who know how to do this. And our role is saying, "We have an interest in seeing this information exist and we can provide some funding so you can go out and do this work."

36:53 CH: Yeah.

36:54 S7: So, you've recommended some charities, especially the Against Malaria Foundation for some years. Have you seen any evidence that the cost effectiveness has diminished as we picked the low-hanging fruit or saturated those charities with funding?

37:06 CH: Yeah. Have we seen evidence that the Against Malaria Foundation or really any of the organizations we recommend that their cost effectiveness has gone down over the years? I don't think that we have seen strong evidence that that's happening other than... Well, we're seeing low hanging funding gaps being filled, so I think that we... Our expectation is that charities will be working in areas that might be less cost effective over time because they are more challenging to work in or because their need isn't the same as the organization where they are currently working. And I believe that's been the case for AMF going into some new countries that we think are less cost effective this year. Is that right, Elie?

37:42 EH: Yeah. Our cost effectiveness have generally... Estimates have generally degraded over time, though it's mostly because we've created... We've essentially thought of more reasons that we should discount the effect that we initially thought was straightforward. But that's largely internal with GiveWell drive cause rather than what we see in the world. And then, in terms of what we see in the world, I don't think we know enough to say whether or not the low hanging fruit is going away and the cost effectiveness is getting worse.

38:20 CH: Okay.

38:25 S8: What do you optimize it for when you say like you said, "Oh. We wanted to have this. We're most interested in long term benefit or the person [38:34] ____." So overall, that leads to your best range. So, what are you optimizing for, staying alive, or having a good life when you're older? Or, what is the metric?

38:50 CH: So, what is GiveWell optimizing for? In the case of deworming, we're looking at long term consumption, as Elie mentioned. So, how much is a person able to spend. With the Against Malaria Foundation and Malaria Consortium,

we're looking at averting a death, typically of a child under five. And so, in order to compare those two things, we end up making a subjective conversion. So, we are looking at different outcomes there and putting in a factor where we're saying, "We think that averting the death of a child under five is X comparable to increasing income in the long run for someone and that that requires a significant amount of both subjective judgement and the data that we do have available on sort of how people viewed having different diseases." And one of the projects where Elie mentioned this group called IDinsight that we're working with on the GiveWell Incubation Grants program. And one of the possible projects we've talked about with them is doing a survey around beneficiary preferences to better understand how people who are actually receiving the interventions make these types of trade offs, since it is a challenging comparison to make when you're looking at these types of different outcomes.

40:04 S9: So, as you start to move from being a charity evaluator towards also being kind of a charity incubator to help the charities you see in the world, and you start to immediately see these charities that are sort of designed and founded with the goal of becoming GiveWell top charities. What are you doing to sort of avoid any notion of regulatory capture or bias? These people who are starting charities are going to be your friends by the time the charities kind of start to scale. And how do you guys think about that?

40:31 CH: Yep. How do we think about avoiding regulatory capture or bias in terms of GiveWell Incubation Grants charities or when someone says, "I want to be a GiveWell top charity?" Do you want to talk about that?

40:42 EH: Yeah. I don't think I have any obvious slam dunk answer that is... But a few things. One, it is absolutely crucial to GiveWell's long term success that we, as well as we can, unbiasedly evaluate these organizations. And we have been faced with this sort of challenge many times in the past. We used to recommend PSI, they were our first top charity. And then, we longer thought they were as effective but we had gotten to know them, and we took them down. And then we recommended VillageReach and we visit them, we knew them really well, but we didn't think they were, we took them down. AMF at one point we thought, didn't need more funding, and that was a very difficult decision but we took them down. And so, it's actually just a situation we've been in many times before. But, yeah.

41:32 S9: As a follow up, do you intend to have the people kind of running the evaluation program, the people running the incubation program be actually the same people?

41:38 EH: So, are they going to be the same people? Yeah. They are in practice the same people. Largely, I am the person responsible for both parts of the program. We're a research staff of eight people. I think that the things that I think we have going for us are a closer in a commitment to doing this well because I actually think the existential risk that GiveWell faces is not doing high quality research. And if we don't do that, I think we will just die, and that is coupled with the commitment to transparency. And so, I think that as long as we have, as long as we continue to publish the full details of our analysis, we will get called out by the people who follow our research if we don't meet that standard.

42:25 EH: Finally, there is an example where we didn't recommend an organization. So, New Incentives I mentioned earlier, they're doing a conditional cash transfer program in Nigeria. This is actually not the first program that they tried to implement. They initially started running a conditional cash transfer program to encourage facility delivery. And when they came to the end of that program, and we were considering them for a recommendation at the end of last year, so in September and October 2016, we were not convinced that the evidence that incentivizing women to deliver in facilities reduced neonatal and maternal mortality, such that we would recommend them as a top charity. And so we said to them, "We are not going to recommend this, but we think you are an effective organization, you have the right values, and you're trying to do the right things. So, if you want, we are willing to provide additional funding for you to try something else." And so, that's what they're doing now.

43:27 S?: So, what are they doing now?

43:28 EH: They are the ones who are doing the conditional cash transfer program to incentivize access to routine immunization also in Nigeria.

43:40 CH: Yeah.

43:40 S10: So, as you evaluate these charities, and they do show a strong evidence maybe in the beginning then they expand, how do you kind of make sure the external validity stays as they expand to new places? And how often do they have to keep showing that it's working where they are in new places?

43:58 CH: Yeah. So, how do we kind of keep up with external validity and think about that in our analysis of charities as they are growing into new places? That's definitely something that we think about. So, trying to think about where a charity is implementing, what are sort of the baseline conditions if we're looking at maybe prevalence and intensity of worms, if we're talking about a deworming organization and how that might compare to a study that we have of it, if we're looking at our randomized control trial that was done. And Elie mentioned including this in the cost effectiveness analysis and making discounts for things like external validity and replicability when it comes to deworming, that's something we do with all of the charities that we recommend. And so, if a charity were to expand into a new country where we either had extremely limited data, where we really didn't have a strong picture of what we thought the intervention was likely to be there, or what the conditions were likely to be there relative to where we had seen it in the past, we would want to discount for that, or understand what the differences were such that it might affect our bottom line cost effectiveness of the charity working in a new area.

45:04 EH: Yeah. And I think it's also going to just depend a lot on how they expand and where they expand to. And this is not a problem we've faced yet 'cause the organizations are so new that they've largely operated in one geography. So, New Incentives case, we keep talking about them. We're talking to IDinsight about running a randomized control trial for their program in Northwest Nigeria, where they're going to operate. Northwest Nigeria has very low immunization rates and a very large population. So there's a lot of room for them to scale up mostly in that context. Now, if they were to move to southeast, some other region where immunization rates are much higher or the culture is sufficiently different, we'd have to think about what sort of additional evaluation to perform to have confidence that the results would translate.

46:00 S11: Can you talk a little more about your outreach strategy? I think that's a really good thing for you guys to do because I think there's positive sort of effects on non-profits that you will never recommend, and on the sector overall, just from raising the bar sort of for everybody. And yet you're kind of like a well kept secret in some ways. And I think that's unfortunate given the power of your methodology and the fact that you're constantly looking and making it better. You're not ashamed to put your failures out there. I think it's incredible. But again, can you talk about what you're thinking 'cause I think it's a great idea [46:35] ____.

46:37 CH: Yeah. So, can we talk about our outreach strategy? So as I mentioned, I might not have a very satisfactory answer at this point since we're early in thinking about what we might do. Historically, we have done very little, we've mostly focused on communicating with individuals who already know about GiveWell and have questions or might be interested in updates on what we're working on. Going forward, we could, as we talked about earlier, sort of target finding like-minded individuals. So, we know some things about our current group of donors, we know that many of our donors work in finance and in tech, so maybe there's something that we could do to reach more people in those sectors. We could also focus on doing things like trying to get GiveWell in the media more. We've been really fortunate at having had some good media coverage in the past and a lot of donors have told us that they've heard about GiveWell from the New York Times, or the Wall Street Journal. And so, should we be more proactively think about coverage for GiveWell in the media.

47:40 CH: Another potential type of area, we could really focus on just networking within Silicon Valley, since we are here and we know many of our donors work in Silicon valley, and that might be another place that we could go. Or doing something more mass marketing-y, but I think that's the least likely of all of the options that I've mentioned. Or even things like improving GiveWell's SEO, improving our website. I think there are also a few sort of things along those lines that might make sense to do, and we're sort of thinking about how to prioritize all this.

48:13 S11: Let me say something. Well, I can already sort of picture you guys jumping on it 'cause it's different when you look at it. But have you ever thought about looking at causes that mean a lot to people? And I don't mean bending your methodology to the point where you can break it but one where you would have causes and issues that were maybe closer to issues that are going to get you press and attention. People that have had worms, it's bad. I was listening to you talk and I'm thinking, "You ever had worms?" Forget about long term just like, it's not a fun thing. But most Americans don't get up thinking about those things. They just don't. Contraception maybe, but not so much worms, and not so much bed nets. And so, have you ever thought about trying to find... Well, I would say find a winner. I know how it goes, it's antithetical the way you think. But if you had an issue, if you had an organization that you could champion with your big brains and your full hearts but that could get you more sort of traction. Have you thought... Again, that's a really kind of sideway way to look at what you're doing, and it's antithetical to the way you view things. But if you do both at the same time... [chuckle] That would be just a win-win, right?

49:30 CH: Yeah. So, could we kind of make GiveWell recommendation for something that we think would just get a lot of PR? I think it's the last point that you made where, can you do both?

49:40 S11: 'Cause even more than PR though, it's you engage somebody where they're coming to... You're hitting them, you're trying to not talk them into like, "You need to understand worms, or you need to understand bed nets," but they already understand. I don't know. Global warming, family planning, whatever it is.

49:56 CH: Yeah. But I think trying to do both simultaneously is an impossible task. GiveWell's brand is that we really care about evidence and that we have this very strict criteria.

50:06 S11: I know.

50:06 CH: And so, just adding something into the mix where we say, "But we make an exception here." Now, there is one exception in my mind that we have made in the past a little bit where we've made kind of squishy recommendations around disaster relief since that's in the humanitarian space, and we do get a lot of questions on that...

50:23 S11: So...

[laughter]

50:24 CH: But we've made those recommendations...

50:26 EH: I think disaster relief is a good example of what you're describing. It's like in the news when something happens... I don't know if anyone does this anymore but you could see on CNN for two weeks straight, pictures. And anytime when donors are really looking for recommendations about where to give. So, I think that's an example of a place where we've tried it. It's a difficult, for the reasons that you describe, it's a difficult fit for what we did. And I think in short... I'm sorry. For what we do, for our methodology, I think our outreach strategy is largely focused on dollars, dollars to our recommendations as the primary outcome metric. That's what we're targeting. And there, I think is just an extraordinary amount of opportunity that we haven't tried to take yet, to reach people who are like-minded, for whom the product we have is what they're looking for. Now, that doesn't mean that we shouldn't ever go beyond that. I've literally sat at events like this for years where people have said, "Well, when are you going to do outreach?" 'Cause it always comes up. And we say, "We're prioritizing research over outreach. Sorry, we're not going to get to that. It's been... Okay." And now, we're turning the corner and we want do more in outreach.

51:36 EH: But I think the first step is going to be to try and just go after the really low-hanging fruit because we've done so little historically, that I really think it will be there.

51:45 CH: And this will probably be... We have time for one last question then we should take a break. So, I want to make sure we get to someone who hasn't had one yet.

51:52 S12: Sorry. You can also dismiss it 'cause it's sort of similar to the previous question. But I was wondering if there's a way to an opportunity to filter domestically since I think a lot of people currently are like... If people who are for instance worried about the current, they're worried about climate change or environmentalism and want to do something just a little bit closer to home. Is there an opportunity to say, "This is the most effective climate change organization [52:17] ____ within the US [52:18] ____."

52:20 CH: Yeah. So similar question. Should we have GiveWell for US? And I guess, the same considerations would apply there in terms of thinking about the challenges of doing that and maintaining our brand of really focusing on the most cost effective opportunities. And also, to Elie's point about that we do believe there is most likely low-hanging fruit around people who haven't heard of GiveWell but are already probably excited to give in this way, and that we should first try to reach those people.

52:51 EH: Sure. So this is, I think an interesting question because GiveWell's initial vision, if you look and you can, if you go on the website, our initial business plan, this was the goal. This is what we envisioned was having recommendations in all these different causes that donors would be interested in. And in 2010, we actually had something of a menu like this. We recommended a handful of domestic organizations and we had recommendations for HIV/AIDS and for family planning, I think. And what we saw empirically is that most of the donors who came to our website and were really interested in our work. Our strongest fans just gave to our top charity at the time. And we saw

that this menu, the sort of menu approach, got very limited uptake. And there's a blog post, if you want to go find it called "A Big Picture Shift in Priorities" from January 2011 or something, where we described why we're shifting from the menu approach to the really targeted approach.

54:00 EH: And so look, that is six and a half years ago. Although, the world has changed, we've changed. We're much better known today than we were then. And so, there's lots of ways in which our experience then may not be representative of what we've experienced today. But one of the things that I've often found just in talking to donors is that the donor in the middle who says, "I really want you to tell me where to give in this cause," and they're totally open minded within a certain area, it's just a very rare type of person. And what we have found much more, and certainly among the people who give the most is that they tend to be extremely open minded to the point where they really want to hear from us, what our top recommendation is. And so, I think if we were going to go down this path, one of the main questions we tried to ask first is whether there's really demand for that sort of product because when we had it in the past, we saw very limited interest in it relative to the GiveWell approach you see today.

55:05 CH: Alright. Well, it's 7:55 so we should break here. We're going to take about a 10 minute break and then Holden will give his presentation on the Open Philanthropy Project. So, be back here in 10 minutes.

[applause]

<https://scribie.com/files/8725cf2f4e9346d0bc947b42731dbc62eac85ac9>