

GiveWell SF Research Event June 19, 2018 – Top Charities

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

00:00 Elie Hassenfeld: Thanks, everyone, for coming. We really appreciate your coming out to the GiveWell and Open Philanthropy Research event. I'm Elie Hassenfeld. I'm GiveWell's co-Founder and executive director.

00:11 Catherine Hollander: Great. And I'm Catherine Hollander and I'm a research analyst that focuses on outreach here.

00:16 EH: Alright. The basic plan for what we're going to do tonight, Catherine and I are going to talk a little bit about some news on the GiveWell front, then we're going to take questions. We'll break up and have a bunch of time for additional mingling and conversation. You can talk to each other, ask us questions. We'll be sticking around and other GiveWell staff will as well, and then reconvene back at 8:15 for Holden's presentation on Open Phil. Just a quick heads up, we are recording this and we're planning to post the recording on our website. That's something that you might have heard that we're into is transparency. If you ask any questions that you want us to take out of the recording, we're totally happy to do that, so just come and find us after, or shoot us an email at info@givewell.org.

01:01 EH: I know there's a sort of a wide array of level of knowledge about GiveWell. I'd like to kick things off with just a little brief intro to what we do. It's probably going to be repetitive for some of you, but hopefully, will set the context and then we'll move through some of the updates from our research this year. GiveWell's focus is trying to find the charities that have the highest impact per dollar, and we largely focus our research, we do focus our research on charities that work in some of the poorest parts of the world. They work on global health and development and do things like try to prevent malaria, treat children for parasitic infections or just give cash directly to very poor people. We select these charities that we recommend using some criteria, which are essentially questions that we're trying to answer to decide whom to recommend and how to rank them. The three basic criteria that we rely on, the first is what's the evidence of effectiveness? How do we know that the program actually works and how well does it work?

02:01 EH: For example, if you're distributing malaria nets, how do we know and do we know that distributing malaria nets prevents malaria and that reduces deaths? If an organization is the one that's actually doing that implementation, are the nets that they distribute, do they actually reach people? Do people use them? Do the nets stay in good condition? And answering that question is really the first big piece of our research. The second one is cost-effectiveness, trying to determine how much it costs the charity to have the impact that it's aiming to. So, how much does it cost to avert a death from malaria? Or how much does it cost to increase someone's income by a certain amount? And then the final criterion that we focus on is room for more funding. Presumably, many of these charities already have significant sources of funds. And so, the question that we're trying to answer is how will they utilize additional funds that they could receive as a result of GiveWell's research?

03:00 EH: And we use these criteria to create what's a short list of Top Charities that we put on our website, that we're trying to direct donations to. In 2017, we directed about \$117 million to those charities. And really every year in our research process, we're trying to answer the questions that could have the biggest impact on how we try to direct those donations at the end of the year. A lot of the donations that come via our research are given in December, near the end of the year, and so we're trying to have the best information available at that time. What we want to go through tonight is basically three types of questions that we ask to try and improve the recommendations that we make. The first is new research and updates on our existing Top Charities. These are groups that we've followed for a long time, we think they're great, and we're always trying to get additional information that could change the way in which we rank them, or it could lead us to stop recommending a charity that we've recommended in the past. Catherine's going to talk about that in a minute.

04:09 EH: The second type of question that we're asking is, are there new charities out there that we could identify that would be competitive with or better than the charities we already recommend? So I'll talk a little bit about the work we're doing there. And then finally, we also want to help provide support for the development of what could be future Top Charities, and we call this GiveWell Incubation Grants. And the idea here is that often organizations don't have all of the support they need to build the monitoring systems that we would look to for the data we would need to do our research, and so if an organization is really promising, but they just don't have the capacity to do the monitoring that we would need, we can provide them some funding to enable them to reach that level of being a GiveWell Top Charity. And there's other types of support that we provide to try and support the development of those future GiveWell Top

Charities.

05:03 EH: Overall, we're just going to be able to highlight a few of the things that we're working on, and so in the Q&A or in the more general session after, I'd really encourage you to ask about whatever is on your mind. We're happy to talk about other things as well, but now I'll turn it over to Catherine to talk about some of the updates on our current work.

05:22 CH: Yeah. I'm going to talk about three different core questions that we're trying to answer this year to inform our understanding of our current group of Top Charities. We currently recommend nine Top Charities. These questions are targeted to better understanding those. The first I'm going to talk about is some of the new research that's been done on cash transfers, which is the program that's implemented by one of our Top Charities, GiveDirectly. The second I'm going to talk about is the concerns that we have about insecticide resistance, which relates to the work of one of our Top Charities, the Against Malaria Foundation, which distributes insecticide-treated nets to prevent malaria. And then the third is on a broader question, which is: How should we think about the trade-offs between the different outcomes that our Top Charities can have, which is a really important input for us, and some work that we're doing to try to get better information about how the beneficiaries of our Top Charities make those trade-offs.

06:23 CH: So I'm going to start with cash transfers, because a few folks here submitted questions when we sent out that link about pre-submitting questions. There were, I think, four about cash transfers. So this definitely I know is an area of interest for folks who are following our research. In short, there has been some new research that has come out this year that raises some questions about the impact of cash transfers, so this intervention to just give people money to buy the things that they want. And specifically, it raises the possibility of two ways that cash transfers might be worse than GiveWell previously believed. The first is that there might be negative spillovers, or the people who do not receive cash transfers might be worse off. And then the second is that the cash transfers might not have a benefit that extends as long as GiveWell was thinking that it would. So I'm going to talk about each of those and our plans for following up.

07:22 CH: So, currently GiveWell doesn't assume major negative or positive spillovers, so we don't assume that people who don't receive cash transfers are significantly harmed or significantly benefited. There was a new study that came out this year that received some attention in the development econ blogosphere, so if you follow that world you might have heard of this. But it was a three-year follow-up study on GiveDirectly and basically looked at the potential for there to be negative spillovers and found some potentially bad effects for the people who didn't receive cash transfers in this study.

08:01 CH: This study is not resembling GiveDirectly's current program. So, the way that this study was set up was that they were looking at transfers that were given to half of eligible households, which are very poor households within a given village, and the transfer amounts were about \$300. The current GiveDirectly program is to give cash transfers of about \$1,000 to all eligible households in the villages that they work in in Kenya and Uganda.

08:29 CH: And so one of the things that we're waiting on from GiveDirectly is midline results from a large study of 650 households that are in a program that is much more similar to their current program, where they're doing the whole village rather than some eligible households within a village. And so, we're planning to wait to do kind of a deep dive on that question until we have those midline results, which we think will be really important.

08:58 CH: And another point on that is we also just don't have a good understanding of the mechanism through which the negative spillover effects of that size would work. And also, we still see that there is a possibility for negative spillovers, but our guess is that it would happen across villages in the current model rather than within villages. So, we'd want to look at that.

09:17 CH: And then the second is on the potential for cash transfers to not have a long-term benefit, but really for most of the benefits of receiving cash transfers to accrue in the first year after a transfer is received. We are currently modeling the benefits of cash transfers in two ways. One is the immediate benefit on short-term consumption, so your ability to just buy things right away. And then also the potential to invest some of your transfer and have benefits over the longer run, like 1 to 20 years.

09:51 CH: So we are planning to re-examine the duration of benefits that we believe cash transfers have, but we also see a sort of important lower bound on our cost-effectiveness model for cash transfers, by virtue of the fact that we currently estimate that about 60% of the benefits of cash transfers come from a short-term consumption benefit. So even if we believed that there isn't a longer-term benefit, we still see significant short-term benefits to recipients and so there's a possibility that the effectiveness, the cost-effectiveness could drop, but there would be a floor on that in our

current model.

10:26 CH: And so as I mentioned, we're waiting for the midline results from this large RCT to come through that we're expecting in the next couple of months, and then we're planning to do kind of a big relook at the area of cash transfers to refresh our recommendation of that program and our understanding of its cost-effectiveness. And so we're expecting to do that... As Elie mentioned, our recommendations come out at the end of the year, and that's when we aim to have the most up-to-date information, so that's our current timeline for working on that cash transfers question.

10:56 CH: The second topic that I wanted to cover is insecticide resistance. So this is a question that we have, which is how much should we be concerned about the potential for insecticide resistance to be growing and what does that mean for the impact of the insecticide-treated nets that the Against Malaria Foundation distributes? Our impression is that this insecticide resistance might be increasing in Sub-Saharan Africa, but we haven't seen any high quality studies that estimate the impact of insecticide resistance on insecticide-treated nets' effectiveness in the field.

11:32 CH: So we have pretty limited information on this right now, and we have to make our best guess in our cost-effectiveness model for how much we think insecticide resistance should impact the benefits that we're modeling of insecticide-treated nets. And so there's a new type of net that's being tested and studied right now called PBO nets. So it's basically a new type of insecticide-treated net that has a synergist involved, which I'm not a biologist, but my understanding is that it impacts the metabolic process by which mosquitoes will internally remove their ability for insecticides to harm them. So that's a very unscientific explanation of it, but basically, it makes insecticide-treated nets potentially more effective in the future. And so right now there's been one randomized control trial of PBO nets that we've looked at, which provides some evidence that these nets provide additional protection over the standard insecticide-treated nets that we see in areas of moderate insecticide resistance.

12:35 CH: And so, we've adjusted our cost-effectiveness model to account for the fact that the Against Malaria Foundation plans to distribute some of these nets in some of its upcoming distributions. And we see that study as kind of an interesting proxy for the overall question of how much insecticide resistance has developed, because we're comparing the impact of a new type of insecticide-treated net with the old one. And so, you might see that as a reasonable proxy for how effective insecticide-treated nets were when they were introduced relative to the non-insecticide-treated nets.

13:09 CH: So, our next steps are that we're planning to wait for results from another randomized control trial, which is a particular type of trial that we rely heavily on in our research that's been done as part of the Against Malaria Foundation's Ugandan distribution that they did in 2017. So we expect to have a bit more information on the effectiveness of these nets and are waiting on results from that trial and following that space.

13:36 CH: The third area was the question that I mentioned that was a bit broader, which is this trade-off that we have to make between thinking about how we, as donors, or anyone in this room who's choosing where to donate to charity, value outcomes that are different from one another. So if GiveWell is comparing an organization like Deworm the World Initiative, which works on a program that we believe increases incomes, with a program like the Against Malaria Foundation, which we believe averts deaths, in order to estimate their relative cost-effectiveness, we need to have some sort of trade-off in our model for thinking about the value of those different outcomes.

14:18 CH: And so, right now, we have very limited information to inform how we make those judgments, and in particular, a type of information that we have not found, and we have looked for this information, is on how the beneficiaries of the Top Charities would make those same trade-offs. We've looked a bit into research that's been done on moral, we call them moral weights in our cost-effectiveness model, and have basically found pretty limited information overall and very limited information on that particular area. A lot of the moral weights research that is done is done in rich countries, and so we want to create this evidence, because we think it would be very important for us. And so through our Incubation Grants Program, which Elie's going to talk a little bit more about and mentioned very briefly in his introduction, which is a part of our work where we're trying to proactively fund research, and the growth of charities that we think are promising, we've made a grant to a group called IDinsight to work with us on our project to survey the beneficiaries of GiveWell Top Charities or individuals living in areas of the world where our Top Charities work, to ask them about the different ways that they would make these trade-offs.

15:31 CH: And so the status of that project right now is that we've done a pilot in Kenya to try to get a better sense of which methods are most promising for asking these questions and getting information to inform our estimates. And so, we learned through this initial pilot that a lot of techniques don't work for providing useful information to us. So there

were challenges of explaining small probabilities and making that understandable. So just some kind of like survey challenges, and then also one of the questions that this research focused on was, as I mentioned, the value of averting deaths versus increasing income and consumption, and we've found that it's unclear what to make of the results that we got back from that initial pilot that we did, because we used two different framings, or IDinsight used two different framings in that pilot for asking the questions and came back with extremely divergent results with those two different framings. And so we're not quite sure what to make of that, and all of us are happy to talk more about it during the Q&A, but trying to not go on and on forever at this point.

16:37 CH: Right now we're thinking about designing a second pilot, what potential other framings we might want to use, we might be interested in talking to people from a different region. As I mentioned, we worked in a particular area of Kenya for our initial pilot, we also primarily, or IDinsight primarily spoke with women during the initial pilot because they were easier to reach. And so we might be interested in speaking with men as well in the future. So we're in the process of learning from the first pilot, and thinking about the second pilot, 'because we think this is a really important question, and right now our approach is to do our best with the limited information that we have, but we're really hopeful we can get some additional information here.

17:17 CH: And then the very final point I wanted to make, which is different from the sort of big research questions, but I think is really important and was directed by one of the questions that we got in advance of the event, which was, "If GiveWell's top recommended charities are so good why hasn't their room for more funding been used up and sort of what's the overall status of the funding need in the areas that we're working for?" And so earlier this year, we did kind of a deeper dive, looking at the funding gaps, the global funding gaps, so not just GiveWell's Top Charities, but overall, in the world, how much funding need is there for two different interventions to prevent malaria. The first being insecticide-treated nets, which I've talked a good bit about, and then also seasonal malaria chemoprevention, which is the distribution of preventive antimalarial drugs to young children during the peak malaria transmission season, and is the intervention implemented by one our Top Charities, the Malaria Consortium.

18:16 CH: And so we did this project by looking at data from two sources, the RBM Partnership to End Malaria, and data that we got from the Malaria Consortium. And what we wanted to was get an estimate of the total size of the global funding gap for each of these programs between 2018 and 2020. And our best estimate based on this project is that there's a gap of \$640 million for insecticide-treated nets during this period, and the largest absolute gap, just to give you a sense of where this is, is Nigeria, which is \$332 million, and then the DRC, which is \$95 million. So those are two of the bigger gaps that we see [19:00] ____ there.

19:01 CH: And then for seasonal malaria chemoprevention treatments, we think that the total gap is around \$165 million, with Nigeria again having the largest absolute funding gap of \$128 million, followed by Burkina Faso of \$15 million. This data, like all data that we look at, comes with some caveats; there are reasons that you might think that the size of the funding gaps might change over time, that the situation might shift, but the sort of bottom line takeaway that we want to convey here is that we think that there is a significant funding gap that is highly cost-effective for donors to fill in these spaces. So I'm going to turn it back over to Elie to talk a bit about some of the work that we're doing on new Top Charities.

19:45 EH: Yeah, so the first part that I want to talk about is some of the work we're doing to find new Top Charities right now. And a lot of the process that we're using there we call intervention prioritization. And the reason is, there's really two parts of our research process. There's the part of our research that scans the world of potential interventions. An intervention would be something like distributing malaria nets, or distributing de-worming pills. And then there's charities that implement those programs, so many charities could potentially implement the intervention of malaria nets. And one of the things that we found is that the differences in evidence and the differences in cost-effectiveness are much wider between interventions than they are between charities implementing the same intervention, and so we really focused the first part of our efforts on that intervention question of which interventions have the best evidence and are the most cost-effective?

20:40 EH: And so what we have right now is essentially a dashboard that has hundreds of interventions that we'd like to look into to get better information about. And we're trying to prioritize them to look into the ones that seem like they have the most potential to be promising and lead to future GiveWell Top Charities. One example of a program that has just come out of this process very recently is, and I should say that for one minute the way in which we're adding programs into this dashboard is we're following sort of academic literature, so looking at what's coming out in journals, in the New England Journal of Medicine and others, the Lancet. And then also just in conversations with representatives of charities, or academics. We learn about new programs that we want to put into this process. So the example that

came out recently is a program to screen and treat pregnant women for syphilis infection. And the best estimate of the scope of this problem is that about 1.4 million pregnant women are infected with syphilis. If it's left untreated it can lead to complications like a still birth, a neo-natal death, potentially also long-term complications for the child.

21:54 EH: And there's a kind of fairly simple, a fairly cheap point of care, so it can be done right at the spot where the patient is there, to diagnose whether or not she has syphilis, and then also provide a simple antibiotic, penicillin, to treat the infection. As far as we can tell from conversations, it seems like this intervention is fairly neglected. So we talked to the person at the World Health Organization who's I guess in charge of their work on this program, and they told us that they spend less than half their time, one person less than half their time, working on this program.

22:36 EH: We've talked to other funders as well, it doesn't seem like a lot of attention is being paid to this program. It's cheap, and we're talking about a million people, so in the scheme of things our sort of best estimate, just to give you a sense, is maybe it's about a \$10 million funding gap or so, and I think that helps to explain some of why this might be a fairly neglected program. We're talking about a fairly small problem in the scheme of things, but this is a program where we have to do more research, we're going to keep looking into it, but certainly our initial estimates imply that it's as or more cost-effective than some of the charities that we currently recommend. So we're really excited to look into this program more, because it's kind of a... It's not the sort of program that you'd expect to appeal to mass market donors, and it doesn't seem like the existing entities that would normally be trying to address it are taking action. And so we're planning to look more into it.

23:37 EH: Now, part of the problem that we have with this program, is we find this program, we're like, "This is really interesting." And the question is, we'd like to fund it, or we'd like to try and recommend it and direct funds to it, but where do we go? How do we find the entity that would be able to take in money and scale up this program? And this is actually a challenge that has come up a lot in GiveWell's history, where when we start with that intervention side research, sometimes we don't find the charity on the other side that can take in money and scale up the program and, to some extent, that's part of the impetus for this Incubation Grants Program, where one of the things that we can do is provide funding to groups that are trying to get off the ground to implement new programs.

24:22 EH: So one of the grants we're currently thinking about in this category of an Incubation Grant to support the development of future Top Charities, is a grant to a group called Evidence Action. Evidence Action is the parent organization of two of our current Top Charities, Deworm the World, and No Lean Season, they also run a program called Dispensers for Safe Water that does, it's a water quality program, and they have this track record of taking programs that have evidence and then scaling them up and implementing them in a high quality way in the field. One of those programs, No Lean Season, grew out of the Incubation Grants Program. I think the first funding we directed to it was in 2014 and it became a Top Charity at the end of last year.

25:06 EH: And so what we'd like to do, what we're considering doing, is trying to give them funding to scale up the work they do in this domain, and they're calling it Evidence Action Beta, and the idea is to go from what they currently have been doing, which is pick programs on a one-off basis, and then potentially move them up slowly, instead allow them to broaden their portfolio and look at more ideas at the beginning, and potentially lead to more Top Charities over time. And I think this program of syphilis screen and treat is one potential program that could be part of this process with Evidence Action Beta. Overall, they're a group that were... Anyhow, we can talk more about them, but we've worked with them closely over time, and so we think they are really well positioned to fill this gap between research and then taking programs all the way to implementation.

26:04 EH: So that was a lot on our end and we're going to pause here and open things up to questions from you, happy to talk about and answer questions about the things we talked about, but also very happy to answer any other questions that are on your mind now or in the general session that's coming after. Yeah.

26:24 S3: I'm curious, how much you guys think about or how possible you think it is for things like distribution networks to be shared across multiple interventions, 'because it seems like if you're intentionally targeting one group of people, and you have multiple kinds of medical interventions, that instead of trying to have some parallel organization that runs a second intervention, you could theoretically try to find people with existing networks. And I'm curious how much you think about that and if that generally doesn't work or if that is a thing that you think is like a promising way to add interventions to these two charities?

[background conversation]

26:56 EH: Well, I think there's a couple of things, first...

27:01 S?: Can you repeat the question?

27:04 EH: Sure. So the question is how do we think about sort of shared logistical capacity to deliver multiple interventions. So instead of deworming and bed nets, you could imagine them both being delivered by the same capacity, therefore saving costs. I think we definitely are looking for that. And if we find that, that potentially reduces costs and then it enables us to find more cost-effective opportunities. In practice, we have not found cases where that shared logistical capacity is achieving significant cost reductions, and some of the reason is that even among the programs we recommend that are often serving, I guess from our perspective, broadly similar populations, they can serve different groups of people.

27:51 EH: So for example, even within a country that Deworm the World serves, in Kenya, there are parts of the country that need deworming medications because there are high levels of parasitic infections and there are parts that don't. And so even there there are these differences between the population that needs the intervention, and then also distributions often have different schedules, so often bed nets are distributed on a mass basis once every three years, and deworming is distributed either once every six months or once every year. And so it's been challenging to find opportunities where we're seeing big cost reductions. Yep.

28:34 S4: You were talking about your surveys of people that attract inherently two different values. Would it be possible to get some sort of real consumption data rather than survey data? It seems like if you're trying to compare money to health outcomes, seeing like whether people actually would spend their money on health outcomes, on health goods, would be a great way, but I don't know if [28:51] ____ things, just totally hopeless to do that in these numbers.

[background conversation]

29:01 EH: I'll repeat the question which is, have we looked at sort of actual spending data from people to see what their preferences are, so you could see whether or not people purchase bed nets if they're available to them. Josh is the Senior Research Analyst that's worked the most on this, but I think, correct me if this is wrong, but I think that we take it mostly as a given that, based on existing research, that people are not going to purchase the health goods when they have the opportunity to do so. And so their preference for what they... They would choose the money that they have over purchasing the health good. Now, that raises an interesting question about whether or not it makes sense to give them that health good. And in some cases there's a bit of paternalism in the decision to recommend the health good over cash.

29:50 EH: There's also positive externalities to giving the health good, so I think the estimate is you cover 80% of households with malaria nets, then that reduces malaria for the other 20% that aren't benefited, and so there are reasons, even if an individual wouldn't prefer that, that it's better for the entire community, but this is the reason that in the surveys that we've asked IDinsight to conduct, they have focused on preferences for the outcome that we specified, as opposed to just the preference for the... Whether or not they would purchase the actual good. Josh, anything you want to add?

30:25 S5: Yes, IDinsight did do a little bit of work to try to think about whether people are willing to purchase health goods, but they tended to anchor on the market price of the health good in the area where they were working. And so, yeah, there is just this really consistent result that's seen in development economics that people are not excited about paying for preventive health goods. And it just leads to this paradox which is, it seems like people do value preventing deaths and having good health, but aren't going to pay for it, and so there's just this challenge of what do you want to do with that information?

30:55 S6: I see from your online materials that you've been talking with JPALS, the government partnership, and also some other agricultural NGOs... And I guess it's like, generally people are doing development interventions that are not in the traditional health sector that you focus on. Is that because there are too many [31:16] ____ restrictions?

31:16 EH: Yeah, so the question is, we've been talking to groups that sort of go outside the scope of what you would maybe expect by looking at GiveWell's current Top Charities. So GiveWell's current Top Charities mostly are delivering fairly simple interventions at mass and we've now been talking to groups that work on advocacy for tobacco taxation in developing countries, a group that Cindy's mentioning, the Poverty Action Lab at MIT, that is trying to, generally speaking, improve the use of evidence in developing country governments. Broadly, this is part of our

ambition going forward. There's, I think, a lot of historical reasons why GiveWell has focused on the types of programs that we have to date. In particular, like when GiveWell started, Holden, who's my co-founder, and I, we didn't know anything about development or charity and so we really focused on the programs that had the most measurable outcomes, where just as donors, we were able to look at the data and assess whether or not the programs were working. Now, as we've been involved in this longer, I think that we're very interested in looking into some of the harder to measure programs that could potentially have significant effects, because they may offer donors some of the best opportunities to get the most impact with the donations that they're giving.

32:49 S7: How close are we to becoming like an adopted charity of a major organization? And how much do we see that kind of adoption as a sign of validation that we're doing the right thing?

33:01 CH: So the question is, how close are we to being adopted by a major organization? Are you thinking of a major funder?

33:09 S7: So I guess a new funder... Amazon, YouTube, Google.

33:15 CH: If a major corporation used GiveWell, exclusively, how close are we to that? [laughter] We definitely see a long way to go in terms of increasing the outreach that we do and increasing the number of people who use GiveWell, whether that's via one large donor that supports us. I'm not quite sure how Amazon buying a non-profit would work, but we are very focused on this question of getting more growth in our outreach. And so this is one of the other major organizational priorities for 2018 for us, which is thinking... Elie mentioned that we moved around almost \$120 million last year. How do we grow that significantly in the next few years? And you know, I mentioned some of the sizes of the funding gaps that we're seeing, and I think that that has been one of the major reasons that we think that this is a time where we are identifying more high-value funding gaps than we're filling with our current donor base. And so, we've started to take proactive steps toward doing outreach, which has not been a major focus for us in our ten-year history. We hired a Head of Growth who started a week ago, Ben, who's in the back, and we have been working on developing an outreach strategy to try to improve both donor retention and also outreach to new donors to hopefully grow, such that we are moving a lot more money to our recommended charities.

34:41 EH: All the way in the back, yeah.

34:43 S8: Along those lines, what's the biggest push-back or obstacle in terms of... You guys are out there. People know that [34:51] _____ are alive from some of these interventions. And then why don't the dollars flow? \$116 million is a lot, but compared to the [35:00] _____, it's kind of small. And so, what's the biggest push-back, I guess, in the space? Is it lack of awareness or are there other obstacles?

35:13 EH: Yeah, so maybe the people in the back probably heard the question... I mean, we all did, but the question is just what are the biggest obstacles to GiveWell's future growth? If this idea makes so much sense, why aren't more donors following it? So I think certainly awareness would be one major existing obstacle. I don't think that most people... I don't know, maybe your friends, but I know most of my friends, before they met me, they hadn't heard of GiveWell. And so, I think awareness is a major issue, but I still think, as compelling that's the argument for impact-focused giving is to me, and probably to all of you, to most people that's just not how they approach charitable giving, they're giving for reasons that are... I shouldn't say... Are fairly unrelated to trying to have as much impact as possible. Self-expression, what their personal brand is just... Or giving to something that's important to them. So I think those tend to be the biggest obstacles with people that I meet, that they are just not driven by the same goals that GiveWell has in making its recommendations. Yep, right there?

36:27 S9: Have you thought at all about whether... So going off of that question. If you have a group of people that even if you present them all this information, about, okay this is the best way to make the most impact possible, they still want to give to other sets of charities because they're sort of personally interested in those outcomes, or whatever. Have you thought at all about whether there's some benefit to be had by saying, "Okay, well, if you'd like to give within this general context, maybe we can find the best way of making impact in that context, even if it's not the most impact overall," because that's a portion of the market that you wouldn't have otherwise had, at all.

37:03 CH: Yeah, and just repeating the question for the recording, the question is whether we might see some benefit of offering sort of a wider range of recommendations or giving people specific recommendations and causes that they're interested in that fall outside of GiveWell's recommended scope. And so, there's actually a historical context here, where GiveWell started out that way, where we said, "Here's the best US education charity that we can find and here's

the best global health charity that we can find." And we had rankings within different cause areas, and we found that the donors who were using GiveWell's research were primarily just asking us, "Where is the place that I can have the most impact?" and that that was really driving our current group of customers, for what it's worth, and so we ended up shifting toward really focusing on that. And I think my perception is we feel like there still might be a lot of people out there that haven't heard of GiveWell. As Elie mentioned, awareness might not be as high as it could be that might be really driven to give in this way that we think is very impactful. And so I think our starting point is trying to get more awareness of our current approach and that we did try the old way and that it was not what was resonating with the group of people who are using our recommendations, and that's why we have sort of ended where we are.

38:18 EH: Yeah, right here?

38:20 S10: You mentioned a couple of criteria you use to evaluate charities you look at. Are there times when those criteria are in tension and if so, how do you resolve a time... You know, imagine, the charity that has no risk-bound evidence, effectiveness, but it's not quite as cost-effective? Versus the other one, which is vice-versa.

38:40 CH: Yeah, the Top Charities on our list definitely...

38:42 EH: The question is, we have these three criteria but they don't always match up. So sometimes you could have a charity that has really great evidence, but has worse cost-effectiveness. And how do we deal with that?

38:54 CH: Yeah, it's definitely a real situation that we're in where right now, GiveDirectly is the least cost-effective charity that we recommend according to our model, but we think is one of the organizationally strongest and has some of the best monitoring, and there's really good evidence of effectiveness, and so this is an active trade-off that we make today and that we try to weigh when we're conveying our recommendations to donors. Another example is on the other end of the spectrum, Deworming is the most cost-effective thing that we currently recommend, but we have the most questions about the evidence base for Deworming. So sort of the highest degree of uncertainty, but also the highest cost-effectiveness.

39:29 CH: And so some of the evidence-oriented questions do get baked into our cost-effectiveness estimates. They're sort of not fully separable where we have uncertainty about Deworming that is reflected in the final cost-effectiveness figure that we present, which organizational excellence isn't, but we are trying to weigh all these things once charities have crossed a threshold, and so cost-effectiveness is... All the metrics are important, and I guess charities first enter our pipeline because the program they're implementing has strong evidence of effectiveness behind it. So they sort of clear that hurdle first, and then cost-effectiveness is really the second sort of hurdle that we're having charities clear, which is that we're now trying to look for things that we think are a bit better than cash transfers, like a couple of times better than cash transfers, because we think that GiveDirectly is a really good baseline for us to use, since there's sort of an intuitive case that giving someone something other than cash should be better than just giving them cash.

40:28 CH: And also that we think GiveDirectly is an excellent organization with significant room to scale and absorb funding, so there might be a few hundreds of millions of dollars of funding gap there. And so we're looking for things are more cost-effective than GiveDirectly. And so GiveWell has to come to a bottom line conclusion for our donors. At the end of the year, we say, "We think that this is the place that you should give if you fully align with our recommendations and want to maximize your impact according to GiveWell." And so we publish a... Last year, it was giving 70% to the Against Malaria Foundation and 30% to the Schistosomiasis Control Initiative. But if you're a donor that feels really strongly that, I don't know, monitoring is the thing that drives you more than anything else or that maybe you really wait for sort of the clearest causal story, you might want to give somewhere other than those two. You might want to give to GiveDirectly, let's say, for the causal story, sort of clearest. You know, cash goes in, people receive cash, that's the outcome. And so, yeah, we try to just communicate about all these differences among our Top Charities. Come up with a bottom line according to GiveWell's staff values, and publish as much as we can to help you make that decision for yourself as well.

41:33 EH: And I want to say like conceptually, what we're trying to do, though this is overly simplified, is say the evidence criterion gives us some probability that an intervention is having some level of an effect, that has some cost, multiply those two things together, and all of those two numbers are subject to where the charity is intending to go next. So in India, Evidence Action or Deworm the World in India it gets, is actually leveraging a large amount of government support to implement the program, so that what they do there is far cheaper than where they operate in other parts of the world. And so therefore, for us the cost-effectiveness of that part of their move for more funding is better than the rest. And so we're essentially just... This is not like literally what happens when we make the decision, but conceptually

we're aiming for maximizing expected value by just combining the criteria together in that way. So I think we have time for one more question. And then, so right here.

42:37 S10: So I have volunteered at the Center for High Impact Philanthropy at Penn and one thing that was... We had a very similar model that we were trying to assess charities on, but one philosophical difference, I think, was that they were thinking about opportunity costs and tended to stray away from things that had high cost inputs. And so for an example, it would be, doctors were required. So one example is, I think it's the Sightsavers charity, we evaluated it and also really decided not to recommend it because we felt that the doctors were probably using their time on positive things anyways and that the charity was taking away their time towards those activities. I would love to just hear from you philosophically, how you do think about that and then how you do evaluate opportunity costs when coming up with your recommendations.

43:23 EH: Yeah, so the question is how do we think about opportunity costs. The term we've often used is negative or offsetting impacts of a charitable program. So every Top Charity review that we publish has a section where we try to go through these and explain what we think might be happening that could offset the impact that we measure. And we had a very similar experience with some of the organizations that are supporting surgeries in poor countries. So I remember talking to someone who was telling me about a program to perform, that they supported, to perform cleft palate surgeries in India and was telling me about how the surgeon was working around the clock to perform all these surgeries and it was clear that there was just no way that you could get more of that surgeon's time. Maybe you could allocate it from something else, but you couldn't get more of it. And so we very much try to take this into account in our recommendations. In practice, most of the organizations that we currently recommend, or all of the organization we currently recommend, are mostly relying on fairly entry level labor, but that's an explicit question that we try to ask in every case. And I think in the doctor surgery case, is one of the places where it becomes most salient.

44:45 EH: Okay, so we're going to wrap up the formal presentation and Q&A. You're welcome to get up, stretch your legs, get more food. And then, Catherine and I and other GiveWell staff are going to be sticking around, so we're happy to chat and answer any additional questions that you have.

45:00 CH: Yeah, and the break will be 30 minutes long. We heard from feedback on our... Or I think if the time is what, it was, yup. We heard from feedback on our last couple events that there was interest in having more time to hang out and talk to one another and talk to us. So Holden will be here at 8:15 for the Open Philanthropy project presentation. But definitely encourage you to talk to us, talk to others that are here, GiveWell staff generally have GiveWell written on their name tags, so we should be hopefully easy to find, mingling is encouraged.

45:31 EH: Yeah, thank you.

45:32 CH: Thank you.