

## GiveWell NYC Research Event December 11, 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.*

00:10 Catherine Hollander: I am Catherine Hollander. I'm a research analyst focused on outreach at GiveWell. And just for a quick overview of the event, I'm going to be giving an update on what GiveWell has been working on and answering questions for the next about 50 minutes, and then we'll take a quick break and you can get some more pizza. And then we'll come back around 8:00 and Holden Karnofsky, the executive director of the Open Philanthropy Project, will speak about what they've been working on, and then we'll wrap up a little bit before 9:00. I'm going to talk about GiveWell's updated list of top charities, which we published a couple of weeks ago, our research that we're working on, and some of our plans for research into 2018, and then some of the work that we're doing to improve GiveWell's outreach going forward.

01:02 CH: A couple just housekeeping notes. We're planning to record the talk, so I'm going to start recording it now. If you ask a question that you would like to not have included in the recording, just email [info@GiveWell.org](mailto:info@GiveWell.org) in the next 24 hours, and we'd be happy to remove it from the recording. We share the recordings on our website as well as the transcripts, so that people who aren't based in New York can hear the event as well. And then finally I'm going to try to remember to repeat questions for the sake of the recording, so when we get to the Q&A you might ask a question and then I will say it again, and then it's written. That's just the quick housekeeping notes, so I'm going to start with just a very high level overview of our research. I'm going to provide a little bit of background about how GiveWell does research, and what our top charities are. In all cases, more information than I'm going to share, probably more information than you're going to want to read is available on our website, and I would just really encourage you to ask any questions, even if you think they might be basic or they might be covered somewhere on the website, since I'm really happy to get into more detail on anything that anyone's interested in hearing about.

02:19 CH: That all said, just at a very high level, one of the updates for GiveWell is that we think our research is in really good shape right now. We invested in more new interventions and charities in 2017 than we had in previous years, and also continued to expand our GiveWell incubation grants program, which is something I'll speak a bit later about. We're very excited about the shape of our research, and that has led us to some new priorities in outreach, which I'll also talk about a little bit later. GiveWell's goal, for those of you who might be newer to GiveWell, is to find and recommend a short list of top charities that we feel represent outstanding giving opportunities for donors who rely on our research. Anyone can come to [GiveWell.org](http://GiveWell.org) and see our list of top recommended charities. These groups meet our four criteria of being evidence-backed, cost-effective, transparent, and in need of additional funding. And so we do this research to focus on how much of an impact the charities have, in terms of the lives they save or improved per dollar donated.

03:32 CH: And so we're really focused on this core question of which charities do we feel represent the best ways to save or improve lives per dollar, and not making a broad evaluation of many charities, but really honing in on the ones that we feel meet our criteria. This has led us to now have nine top charities. They all work internationally in global health and development. Two new additions to our list this year were Helen Keller International's Vitamin A Supplementation Program, and Evidence Action's No Lean Season Program. And then we continued to recommend the seven charities that were on our list in 2016. So just quickly, for the overview of the shape of our list, three focus on health interventions of our recommended charities. That's the Against Malaria Foundation, Malaria Consortium's Seasonal Malaria Chemoprevention Program. Both of those, as you might be able to deduce, work on programs to prevent malaria, and we think the main benefit is reducing child mortality from malaria. They have different programs that they're doing, but that is the end goal. And then Helen Keller International's Vitamin A Supplementation Program is the third health-related charity that we recommend.

04:50 CH: We also have four charities on our list, as we did last year, that work on programs to treat intestinal parasites, or deworming programs. This is a health intervention where you're treating intestinal parasites, but the main benefit that we see and the reason that we recommend it is due to deworming leading to a potential increase in incomes later in life for children who receive deworming treatments. Four of our top charities work on deworming programs, and then two of our top charities aim to reduce poverty and increase incomes. That's Evidence Action's No Lean Season Program, and I'll talk more about that one since it's a new one, and then GiveDirectly's work to distribute direct cash transfers to very poor households. They just give money directly, with no strings attached, to very poor households. Details about all these groups are on our website, but again, I'm happy to talk through updates or describe a bit more about what any

of them do if you're interested. And then just quickly running through how we get to this list, how we find organizations in the universe of many potential things that we could evaluate, how we decide to evaluate these, why are they on our list.

06:02 CH: We start by reviewing the academic literature for programs that seem promising. This is sort of independent of looking at particular charities. The first step in our process is saying, "Is this type of program likely to be strongly evidence backed and cost-effective?" The reason that we start there is because, well, there's two reasons. One is that most charities aren't collecting that type of information about their work, so they're not running randomized control trials, which is a type of academic study that we find very helpful in our research. They are usually focused on implementing the program that they're implementing. The Against Malaria Foundation isn't focused on proving that receiving an insecticide-treated net reduces child mortality from malaria, they're focused on getting nets to children and then the evidence that they provide is that they were successful... Or getting nets to people, and then they demonstrate that people received those nets. So we start with that. That's one reason that we started the program: Evidence.

07:05 CH: The other is that the major differences in cost-effectiveness that we see among different things or different charities are really often driven quite substantially by the programs that they're implementing. It's a very good starting place for us for that reason since cost-effectiveness is really important to us in our final recommendations. What we do throughout the year is keeping an eye on journal articles for new studies that are coming out, talking to experts in the space, and just keeping an eye on what's happening in the global health and development space to see whether there are programs that we might want to assess to potentially identify charities that are working on those programs.

07:47 CH: Once we find a program that we think has this body of evidence and appears to be strong and cost-effective, then we'll start looking for a charity that's implementing that program, and that was the story of Helen Keller International's Vitamin A Supplementation Program, which is new on our list. We started by saying, we would be really excited to support a vitamin A charity because we've been looking at the evidence there and it seems really strong, but we didn't have a charity yet in mind, and eventually got to Helen Keller International's Vitamin A Program.

08:19 CH: That kind of suggests the next step in our process, which is reaching out to organizations that are implementing promising programs that we've done that initial review on, and we have a staff member now who's focused on doing that type of outreach. This is a newer development for us in 2017, which is trying to better communicate with charities that are implementing programs to say, "You're doing something that we're excited about, you should consider applying for a GiveWell recommendation. Here's what we'll do it for, here's the type of evidence we are interested in, and here's what our recommendation means for organizations that receive them." That's how we ultimately winnow down the types of groups that we're looking into. And then, as I mentioned, we're not trying to be a comprehensive evaluator. Our process kind of funnels the amount of time we're spending into really focusing on spending the most amount of time with the charities that seem the most likely to become GiveWell top charities.

09:15 CH: Once a charity seems unlikely to meet our criteria, we will de-prioritize work on that charity and really try to focus on spending the most time doing the deepest dives into the groups that seem likely to be on our list. We have nine top charities. Some of the strengths and weaknesses that they have relative to one another, and are some of the things that we're weighing when we're thinking about whether to recommend charity A or charity B to donors: A core question that we have is how cost-effective is the charity's program. I mentioned that this is one of our key criteria. It's a really big part of what we do, spending a lot of time thinking about the cost-effectiveness of potential top charities and the ones on our list. That being said, we do think that our cost-effectiveness model is subject to significant uncertainty, it involves a number of subjective value judgements and moral value judgements that can vary from individual to individual. We see this as really important, but not the decisive, only factor that we take into account.

10:22 CH: The benchmark that we're using when we're assessing the cost-effectiveness of a potential top charity is the cost-effectiveness of cash transfers. GiveDirectly, the charity that I mentioned that distributes direct cash transfers, we think is an excellent organization with significant capacity to absorb hundreds of millions of dollars of additional donations, and also is implementing a program that provides a nice, intuitive baseline for us to think about comparing programs to. When you're thinking about helping someone, it kind of intuitively makes sense to think about whether that's better than just giving them money to buy the things that they want to buy, and so we are generally looking for things that are better than cash transfers by something in the range of two to three times, with some squishiness there, since we don't think that our cost-effectiveness assessments are precise, as I mentioned.

11:17 CH: And then, some of the other questions that we ask when we're looking at organizations, one is how well do we know the organization. We have now had some of the top charities on our list for six years. We feel like we know

them quite well. We also have newer organizations that we don't know as well, and so when we're thinking about how much funding to direct to a particular group, that's a factor that we're thinking about. We're also interested in comparing the relative strengths and weaknesses of the organization's monitoring capabilities. What I mean by monitoring is the type of work that I mentioned that the Against Malaria Foundation does when they are asking if a person has received an insecticide-treated net and if it's hung up after the fact that they've gotten one; or after deworming pills are given out, checking to make sure that the kids received the deworming treatments as expected.

12:09 CH: All of our top charities have strong monitoring relative to our understanding of the average charity, but they do vary in their strength of monitoring among the top charities. That's one factor that we're looking at when we're comparing our top charities to one another as well. And then finally, just the ease of us communicating with the charities is an important factor for us, since we have a lot of questions and some charities, it's a little easier to communicate with than others and that's kind of a final factor in terms of us thinking about how to compare the strengths and weaknesses of our own top charities. That is a factor when we think about what to recommend to donors. When we're thinking about all these things put together, we come up with a bottom line recommendation to donors, saying, "We have nine top charities. If you want to follow GiveWell's recommendation, here's our headline recommendation."

13:03 CH: Which this year, is to first give to GiveWell for grants at our discretion, which means that we will grant the funds to the opportunity that we see as the most pressing at the time that the donation is granted out. Or secondly, if you prefer to give directly to a charity, to give 70% to the Against Malaria Foundation, 30% to the Schistosomiasis Control Initiative. Just quickly touching on the new top charities, I mentioned that we added two this year. The first was Evidence Action's No Lean Season program. What they do is that they run a seasonal migration program in rural northern Bangladesh, where they provide no-interest loans to low-income agricultural workers during the time of year in which agricultural job opportunities are low, and it is the most challenging time to find work. It's the period between the planting of the major rice harvest and harvesting. The no-interest loans are conditional upon a member of a household migrating to another area of Bangladesh to find employment opportunities. This is something that we initially came across when they did an academic, Mushfiq Mobarak, who's at Yale, he's an economist, and ran a randomized control trial, which is this kind of gold standard type of study that we value very highly in our research. He ran a study of this program in 2008, and we came across that study and thought it...

*Due to a recording error, the below portion of the research event was recorded separately from the above portion. This recording can be found on our GiveWell Public Research Discussions page under the "December 11, 2017" event labeled "Audio recording of the event – Top Charities Part 2." There is also a small gap in the audio between the first and second recordings, reflected in this transcript. 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 CH: We actually need an initial, what we call an incubation grant, which is this separate piece of our work where we're trying to help build potential future top charities and to improve our understanding of our current top charities. So we made a GiveWell Incubation Grant to No Lean Season back in 2014, since we didn't think at that time that they were quite evidence-backed enough that we were comfortable recommending them as a top charity since we want to see a lot of evidence, as I mentioned, we have a pretty strict criteria. We're really looking for organizations with a very strong case of impact. And No Lean Season was very new. But we provided this incubation grant back in 2014 and followed their work for the next few years until we got to 2017, when we felt that they were operating at a sufficient scale and had the type of track record that we could assess for a potential top charity recommendation.

00:50 CH: So this year we spent a lot of time asking them questions, really modeling out their cost-effectiveness, figuring out how much additional funding they need, and then also doing a site visit to Bangladesh, which I was fortunate to join. So checking out the program on the ground, watching how it works, and asking lots of questions, basically spending a week asking Evidence Action and their implementing partner, RDRS Bangladesh, as many questions we could possibly make up. [chuckle] So that was how No Lean Season joined our list.

01:25 CH: And then Helen Keller International's Vitamin A Supplementation Program, I kind of quickly hinted at this earlier. They were implementing a program that we're very interested in for a while, Vitamin A Supplementation, and Chelsea Tabart, who's the research analyst who's been leading our charity outreach work, reached out to them to potentially apply with this program. And so we spent the year similarly, really trying to understand the program, understand the evidence base for Vitamin A Supplementation better and modeling out the cost-effectiveness. And then

finally, we believe that there was a funding need this year in particular because the Canadian government had scaled down funding that it had been providing for this program. So we saw that additional dollars could really have an impact here. So Helen Keller International, we also did a site visit there to their work in Guinea to observe a vitamin A supplementation distribution, and that's how that organization also joined our list. Those are the two newest additions to the GiveWell group.

02:32 CH: And some of the considerations that we would recommend as donors, keeping in mind, I mentioned that, for one, our cost-effectiveness models involve some uncertainty, and that uncertainty can reflect a few different things; one is the ways that staff members at GiveWell trade off increase in income relative to improving health. As I said, our top charity work on programs with pretty different outcomes, that tradeoff between health and income is a quite subjective one at this time, where staff members will input their own best guesses as to how to trade this off. And we're actually working on a project to get more information about how to make those tradeoffs that I'll talk about in a moment. If you're a donor thinking about giving to one of the nine charities, we recommend and you feel very strongly that improving someone's health is the number one most important thing to you, you might want to choose to support one of the charities that we recommend that focuses on that rather than deworming or No Lean Season, or GiveDirectly, or the opposite may be true. So that's one thing to keep in mind.

03:45 CH: The other major point to keep in mind if you're considering which of the charities to support is your feelings about the certainty of an evidence base that you would want to see. So the top charities that we recommend are all very evidence-backed and we feel that they are all outstanding, but they don't have the same types of evidence basis behind them. So what I mean by this is that you want to really know that your dollar is going from point A to point B, that you might make different choices about which GiveWell top charities to support. And so, sort of on the spectrum of most certain case for impact to most speculative case for impact, we have GiveDirectly on one side which is a program where you're just giving people money. There have been many randomized controlled trials done of this program. GiveDirectly itself has excellent monitoring and in fact, have a very clear case for where those dollars go. Somewhere in the middle, we have the malaria charities, so the Malaria Consortium and Seasonal Malaria Chemoprevention also have been studied quite extensively and academic research. The research might be a little bit older. Maybe Malaria programs are a little bit different where they're working. They also have good monitoring, but the case is slightly less certain, I would say, than the GiveDirectly cash transfer case.

05:05 CH: And then on the farthest side, we have the deworming groups that we recommend where we rely very heavily on one randomized controlled trial that showed a very large impact on incomes for children who received deworming later in their life. We discount this study quite heavily, and I'll probably just talk more about the reasons for that, but it's really the single piece of evidence that drives a lot of our recommendation. And we've spent a lot of time trying to figure out what's going on with this evidence and how it holds up. We do think that it's continued to hold up against our scrutiny, but it's not the same level of certainty as something like the Against Malaria Foundation or GiveDirectly where there are multiple randomized controlled trials that point in the same direction. Our recommendation of deworming is really based on having a very chief intervention with a potentially very large impact. So even when we discount the potential that the impact wasn't as large as in that main study that we rely on by huge amounts, like 90%, it still comes out to be one of the most cost-effective things we've ever seen. So certainty is sort of another thing you might want to have in mind as a donor. And it should say...

06:12 S?: Where does the vitamin A supplement issue follow?

06:15 CH: I'd say vitamin A is probably somewhere between GiveDirectly and the malaria charities. We have some questions about how the vitamin A deficiency levels in the areas where Helen Keller International is operating compared to the areas where the vitamin A studies were done in the past, and also the fact that there's this large number of studies that were done in the '80s and '90s of vitamin A supplementation, and then a very large, more recent study... Sorry. Those studies showed that it reduced child mortality from illnesses than a very large, more recent study that didn't find an effect. And so we've spent a lot of time thinking about that question, so that was a factor in us, figuring out that we did want to recommend vitamin A supplementation and with that. But that would be my best guess, would probably be somewhere between GiveDirectly and the malaria charities.

07:11 CH: And I should say with all of these donor considerations, one thing is that we're really happy to talk through giving decisions, and so if you're thinking about, "What do my preferences mean for what I should support?", or "What does GiveWell think about for donors who feel this way?", we're really happy if you reach out to us to help talk through those decisions or provide more detail and sort of why we recommend what we recommend and which may be values would lead you to also choose direct... Or to donate to that place. So just like I said earlier, plenty of time for questions,

I wanted to talk about two of the elements of GiveWell's future. The first is GiveWell Incubation Grants. So I mentioned that No Lean Season came through GiveWell Incubation Grants that this is a piece of our work where we're actively trying to grow the pipeline of potential future top charities and to improve our understanding of our current top charities.

08:09 CH: And so the thinking behind incubation grants, which started in 2014, was that we weren't seeing a lot of turnover in our top charities list just by going through our standard process of looking for organizations that have a lot of evidence pre-existing behind them and very long, strong track records that we could assess. And so we wondered whether we might be able to get involved as a funder of it earlier to develop organizations that seem promising since the groups that we were assessing had been generally funded by someone else before they got to us such that they could develop the type of track record that GiveWell requires for our recommendation. And so we made a number of grants, these are supported by Good Ventures, a large foundation that we work closely with, as part of our incubation grants program over the years, and we're really excited about this work and we've sort of made it a bigger piece of how we're spending our time.

09:03 CH: Because we feel like we are now able to think about the research questions that are most important to us that could really shape our recommendations and to identify organizations that seem really promising like No Lean Season, and to support them such that we can answer questions that we have or recommend things that we hadn't recommended previously. And one of the ways we're doing this is through a partnership with a group called IDinsight, and IDinsight is a group that provides social impact information on the ground to policy makers and decision makers. What they can do is be a research partner for us that's operating in the field, so they can run a randomized controlled trial if we have a question that we would like to run a randomized controlled trial under, or they can do a survey... In this case, one project that I mentioned is that our cost-effectiveness figures a lot of these moral tradeoffs that are, "How do you think about income versus health?" And that's a really hard question, and some information about this exists in terms of how people make those tradeoffs in rich countries, some or most the research that's been done on how people generally value income and health tradeoffs.

10:13 CH: That exists for wealthy countries and does not generally exist as far as we can tell for the countries in which GiveWell's top charities operate, so some of the poorest parts of the world where you can imagine maybe people make tradeoffs differently. And so one of the projects that we're working on with IDinsight is a survey of beneficiaries of GiveWell top charities or potential beneficiaries of GiveWell top charities to better understand how they would make those tradeoffs. So that's one type of project that we're excited about to be working with IDinsight 'cause that would really help inform our recommendations. Then finally, one of the other major pieces of our work going forward is on GiveWell's outreach work. Historically, GiveWell is 10 years old this year, and we've historically put very little focus on proactive outreach about GiveWell's work.

11:07 CH: We really focus on building up the research product, and we feel like that product is now both quite mature and identifying more funding opportunities than... More high-value funding opportunities that we're expecting to fill with our current projected donor growth rates. And so, we're now thinking a lot about how we can improve our outreach and reach people who might not have heard of GiveWell before but hope will be interested in using our research. And so we've been working on a number of different projects there. Some of the biggest projects... One of the biggest is on an experiment that we're running in podcast advertising, where if you listen to The Weeds or FiveThirtyEight or Rationally Speaking, you might have heard an ad for GiveWell in the last few weeks. That was the result of us running an experiment in February of having a couple ads run on those podcasts and seeing a really good return to that, and thinking that that was something that was worth trying during the giving season, during the time of year when most people are making terrible giving decisions.

12:10 CH: So that's one type of outreach project that we're working on. We're also interested in, it's kind of the marketing category. We're also interested in things like improving GiveWell events, so we'll be sending a survey and it will be a great help if you can pull that off thinking about participants that would be helpful to have, thinking about our media outreach, since lots of GiveWell's donors have found us through media coverage that we've received. And so, generally, just running experiments in the space and trying to figure out what works for us, what doesn't. We're quite new to it, but this has become a major organizational priority, and so we're trying to do more of that going forward as well to learn from the initial work that we're doing in this space, and hopefully, develop a robust outreach plan going forward. So I'm going to pause there. So I've been talking for a while and I would really like to answer questions that you have. So as I mentioned, I'll repeat them for the sake of the recording, and yeah, feel free to ask about any of that or anything else I didn't mention.

13:12 S?: Alright. So it sounds like a requirement for getting attention from GiveWell and getting a recommendation is, your intervention has been studied academically. And so, are you aware of any biases in the kinds of things in academic study that might be neglecting like high impact causes or interventions?

13:34 CH: Yeah, so, is it... Are we potentially sort of having a biased set of research that we're considering that's coming out of academia since that is a requirement for us? So, when I say academically, too, I just mean independently studied. So there are some groups that are either affiliated with academic institutions or separate research organizations that focus on global health and development. So it's not just traditional academia that we're looking at, but we're probably... I think an important question is whether we're missing out on things generally that could be better than our top charities by focusing on what's measurable or what can be studied? And Holden will talk about some of the work that the Open Philanthropy Project is doing, which is kind of in a different category of evidence in some cases where they're open to things that are a bit more speculative than GiveWell.

14:31 CH: So GiveWell is really focused on providing strongly evidence-backed and vettable recommendations, and when I say vettable, I mean we want to be able to make a very strong piece for anyone who sees our research as to why we recommend some things. So even if we don't put [14:48] \_\_\_\_ checked equipments, you can check them and see that there's a strong case there. And so for us, the ability to study and measure something isn't really important prerequisite because we do want to be able to make this strong case toward why we recommend what we recommend. That being said, we are interested ourselves, as I mentioned in the incubation grants program and actively funding research into things that seem promising, but maybe haven't been studied yet. And so we are interested in having conversations to point toward things that we would love more information about, that doesn't currently exist. And so I think that we're trying to generate evidence in some cases where it doesn't exist, but we would like to see it.

15:36 S?: So you mentioned that you're looking for charities that are two to three times more effective than GiveDirectly. Why that two to three multiplier, like how would [15:48] \_\_\_\_ to be okay with things that are at least as good as GiveDirectly?

15:54 CH: Yeah. So why do we look at something that might be two to three times as good as cash or GiveDirectly when we're looking at recommendations? So, it's not a hard and fast rule, as I mentioned. We don't think our cost-effectiveness estimates are precise, and so that's why we look for things that can have 2 to 3x different instead of just 0.25 times as good as cash transfer to something like that. And that's earlier in our process when we're prioritizing among potential interventions. And we often see the cost-effectiveness of things get worse the more time that we spend learning about them and thinking about our uncertainty, and putting in our input, so I would guess that if something is as good as cash transfers in our initial take on it, that will potentially end up being worse than cash transfers, and we want to really solidly be GiveDirectly since we do think that it is an excellent organization with a lot of growth for more funding.

16:50 CH: And so I think that we're trying to aim to really find things that are much better than cash rather than kind of equivalently good because we do think that there is so much room for more funding. If we had a scenario where something ended up being around as good as GiveDirectly, I'm not sure that I could say what exactly we would decide, like if it started out at 10x cash and ended up around the same after doing our deep dive into its cost-effectiveness. So far that hasn't happened. A lot of the interventions that we found this year have ended up like around six times as cost-effective as cash transfer's range, and we're trying to prioritize finding things that might be significantly better than that. Yeah, so I don't mean to be too precise when I say 2 to 3x; that's just kind of like a rough heuristic; we're more in the earlier stages of prioritizing what to look into, to almost kind of gauge how promising something is.

17:43 S?: And it sounds like you expect there to be interventions that are much more effective than that, and so you're looking for those.

17:49 CH: Yeah, we're hoping, so we found a lot of things where the marginal dollar of, here's like it'll be around six times as cost-effective as cash transfers and we think that there's probably a lot of room for more funding in that space, in that range based on what we found recently. And we're trying to think of what would beat that, where we could find opportunities that would be significantly better. Our current best guess is policy advocacy in low and middle income countries, because you could imagine putting in a relatively small amount of dollars, having a policy passed that affects millions of people and it might be extremely cost-effective. So we're planning to prioritize some research into that in 2018. Whether we'll be able to identify opportunities or make the strong enough causal piece that we would recommend that, I'm not sure at this point, but that's kind of our best guess for where things that are significantly better than this kind of cluster of cost-effectiveness that we've found would be.

18:46 S?: Okay, thank you.

18:51 S?: Over what time horizon are you looking at impact? I think about the healthcare stuff and the malaria bed nets and stuff [18:58] \_\_\_\_ obviously, but disease eradication which can be more expensive upfront, but over time can save many lives.

19:06 CH: So what time horizon are we looking at for impact? So for the things that we currently recommend, there is actually quite a variation since the deworming treatments that I mentioned, we think that the main impact is 10 years out. And so, with something like disease eradication, I think that type of question would probably come down more to the evidence piece and how confident we could be that funding a particular thing would lead to a disease being eradicated. I think that would be more of the challenge than something like an expectation that that would occur over a 10-year time horizon. Since we're comfortable with that kind of time horizon, if there's evidence base that suggests that that's where the impact is.

19:50 S?: Hi, I'm very excited by the outreach efforts that you're currently mentioning, that initial round of podcasts advertisements for example, that seem to target people who are inclined to, let's say, evidence-based policy decision-making. Do you have ideas or interests in pursuing how to persuade people whose basis for their charitable giving might not be that framework?

20:18 CH: Yeah, so our initial outreach effort seems to focus on donors who might be data-minded already or potentially already thinking about giving in this way. Do we plan to or have an idea how to reach people beyond that? So, our initial thought is that it makes sense to focus on people like our donors, so we chose the podcast that we chose to advertise on because we talked to a lot of donors about what they were listening to and those were really common options, and so we thought we have a sense that not everyone who might be interested in GiveWell has heard that we exist, and so we should start with people who were most likely to give through GiveWell if they hear about us, and then to kind of broaden from there.

21:01 CH: It seems like a reasonable starting place for us, but I definitely, I'm not sure what we're going to learn out of all these experiments. Maybe we'll find out everyone who's already potentially... Everyone who's the most likely to give to GiveWell is already giving to GiveWell. We don't think that's true, just based on the average that we've already done, but we might [21:17] \_\_\_\_ point that, and then we'll have to probably get into the trickier, "How do you reach branding people who are giving in very different ways?" Yeah, so it's a good question.

21:28 S?: Yeah, you talked about looking into funding for new research into new interventions, and you also spoke about how you wish there was more research for your existing interventions. Have you thought about funding research for those?

21:39 CH: Yes, great question. So have we thought about funding research into things that we already recommend? Deworming is probably the best place where we're really thinking about this, and then there's also Against Malaria Foundation work. So I guess I'll start with that one 'cause it's shorter, and then get into deworming and hopefully be able to explain it in a short amount of time what we're thinking about with that. So with the Against Malaria Foundation, we're funding IDinsight, that group that I mentioned that we work closely with, to work on a project to address some of the questions that we have about their monitoring, which we raised. We wrote a blog post in 2016 where we talked about some of the questions that we realized we had about the way that the Against Malaria Foundation was doing those check-ups that I mentioned, after they do insecticide-treated net distributions to visit villages and visit households that received them and see if they're still hung and what condition they're in.

22:29 CH: So we raised questions about the methodology used by which they were doing that. IDinsight is now working on a project to both better understand how the Against Malaria Foundation is doing that and to work with them to improve that. That's one project. And then on deworming, where I mentioned we rely very heavily on this one study, which we spent a lot of time looking into, I think this is the place where we would probably be the most excited to see additional evidence. And so, we made two additional incubation grants in the deworming space in 2017; one, to fund a 20-year follow-up study of consumption for people who received deworming treatments as children, and then the second, to look into whether there were studies that were done 10 to 15 years ago of deworming where maybe those studies had a one to five-year follow-up, and whether we could revisit recipients of those studies 10 to 15 years later, and see whether they also find this income effect that we see that's really surprising, that it would be interesting to see if other studies also found that effect.



23:36 CH: And so that was another smaller grant [23:38] \_\_\_\_ just to get an initial sense of whether that information could be collected. I think we're pretty unsure about that one going anywhere, but those were two of the things that we are trying to learn more about deworming. David Roodman, who's a senior advisor at the Open Philanthropy Project and specializes in doing replications of studies that they've conducted, recently published a very detailed blog post on the GiveWell blog on his replication of a hookworm eradication study in the American South that informed our recommendation of deworming. He ultimately concluded that that study was not persuasive, so now we have that data point. It's not the core piece of our deworming recommendation and we continue to recommend deworming, knowing the results of David's research. And now, we're in a point where we're thinking about what else we could do. The reason... You might think the obvious thing to do would just be to run another randomized controlled trial starting now. There are a couple issues with that; one is just the cost of doing so. It would have to be a quite large study to have the statistical power that we want to see, and we'd also have to wait 20 or so years to get results [chuckle], and we're directing funding to deworming now, so we would love to fund something or do something in the nearer term that could improve our understanding of deworming.

24:53 CH: And so, we've talked about a variety of options and I think are still weighing internally whether there are other things that we could do that would change our minds, and with all of them, I think our major question is the cost, and then the probability that the information we get from that study would change our minds. And so, thinking about how to square those two things is going to sort of guide whether or not we move forward with any particular thing that we might do to try to understand it better.

25:23 S?: I have to get those... [25:23] \_\_\_\_ sorry here we go.

25:25 S?: Thanks. Most of the recommendations are "programs of a larger organization". Have you all considered the organizational health of the charities that you recommend the specific programs of? And if there are any long-term implications for those organizations of supporting specific programs rather than the overall operations that gives them the income-restricted support, that gives them the flexibility to adapt to changing circumstances?

25:54 CH: Yeah. So the question is, many of the programs that we're recommending are for specific programs that a large charity runs. So we're not recommending the whole charity, we're just saying you should give to the small piece of it. And do we have long-term concerns with the overall, or do we watch the overall health of the organization, such that... Basically, are we concerned that just funneling money to vitamin A supplementation might mean that the overall health of Helen Keller International is changing in some way that's negative? And this is actually kind of a broader point, is this is a newer situation for us where historically, GiveWell hadn't really recommended sub-programs of larger organizations for the most part, out of our concerns with the fungibility of that money with other programs that the organization was doing that we didn't want to recommend.

26:49 CH: And so the fact that, I think five or six of our top charities are now sub-programs of a bigger organization is a shift for us, and reflects the fact that we feel more confident in our ability to assess what additional marginal dollars will do with those groups. So part of our recommendation is thinking about whether the organization will be able to carry out the program in question. So if they don't have sufficient funding for their overhead in such a way that that would negatively impact the programs we're interested in, that should be something that we find out about as part of our asking about whether they have additional funding needs. And then, I'm not sure whether something like a program of a large organization that we recommend... Sorry... We recommend a sub-program of...

27:39 CH: Let's say this is sub-program A and there's sub-program B, and then charity C. If sub-program B isn't what we recommend and had to shrink as a result of funding going toward our recommendation, I mean that shouldn't necessarily what would be happening since we're trying to grow the program of interest, sub-program A, by directing more dollars to it. So if something happened where sub-program B shrunk, I don't think we would, by default, try to grow that program, since it's not something that we recommend. And so I think we would just want to know whether the opportunity that we're most interested in is continuing to grow, and if there were bigger organizational issues that would impact that, that would be something we would want to be aware of. But we don't, I guess by default, look at the status of all of the programs that an organization is implementing beyond how they impact the sub-program that we recommend.

28:42 S?: How do you decide how you allocate dollars between giving money to charities and all versus say, the incubation?

28:48 CH: Yeah. So the incubation grants are fully funded by Good Ventures, the large foundation that we work closely with. And so they are funding that at a particular budget, and then they are also choosing, for their own dollars, how much to support our top charities. And so I think right now we feel like we're able to fund the incubation grants that were [29:11] \_\_\_\_ at sort of the level that we're seeing. In the future, we may have to start making tradeoffs of those dollars, but it's not individual donor dollars that we're trading off; it's this one specific foundation, and so far we've been able to fund those highest-priority incubation grants that we've seen.

29:33 S?: A lot of talk about replications of studies and your recommendation hinging on the statistical power of the study or at least just the... You know that what the study is showing is actually happening on the ground. How deeply into the studies are you able to get? Obviously, replicating studies is extremely expensive. Do you have access to the data of the study? Can you get access to the data? Can you re-analyze that? As a skeptical scientist, I know that I can give you associations between any two variables if you give me enough time to run enough progressions on this data or if I clean it the way I want. I'm just curious if you have a data scientist that kind of does.

30:20 CH: Yeah, so the question is, do we have a data scientist that looks at the data when they're doing a replication, versus I guess us just taking it at its face value?

30:31 S?: Or just even looking over the study that's been done, right? It's been done in this study that claims this 20-year association, if you have access to the data and you can kind of see if you think that the statistics that they claim are actually... If you run it... If you're running the model slightly differently, that the results change. Right?

30:49 CH: Yes. Yep. So that's the type of work that David Roodman is doing that I mentioned, where he's going back in and trying to look at the original data and gathering some additional data himself. And so, I highly recommend if you're interested in the specifics of how he does the replications that he's done, I highly recommend reading his blog posts that came out of how he actually did the replication and what he found. So he is looking at that type of information, as part of his...

31:19 S?: Does he work for you guys or he's just a guest blogger or...

31:23 CH: So he is a, what is it, a senior advisor at the Open Philanthropy Project, and has also worked for GiveWell. So this work that he does is work for GiveWell that he is producing for us to use.

31:38 S?: So a question before we're... You're talking about GiveDirectly versus other programs that maybe work comparably as effective as GiveDirectly, and had brought up from the question of, is the difference between giving directly, which kind of empowers people to make their own decisions, gives them more autonomy versus doing things for them, is that something that plays in as a fact?

32:13 CH: Yeah, so the question is, does the fact that GiveDirectly empowers people or gives them autonomy play into GiveWell's recommendation of GiveDirectly? It certainly plays into a lot of donors' decisions to support GiveDirectly through GiveWell, so it's long been on our list, but it's always been the least cost-effective opportunity that we've recommended, you had mentioned earlier. I think a lot of the donors who do choose to support GiveDirectly, there's kind of this certainty element that I mentioned that appeals to some, but also this empowerment metric. And I think especially for donors who are very concerned about paternalism and paternalistically deciding what to do on someone's behalf, they will choose to support GiveDirectly. So, in terms of whether GiveWell explicitly models that, that's not part of our cost-effectiveness analysis, but I certainly think it's a factor that we mention to donors and that we know that people make decisions based on. So it's sort of one of the qualitative pieces that people use when they're thinking about which charity to support.

33:13 CH: And I think we just have time for one more.

33:17 S?: I was just curious of the organizations that you evaluate, about what percent do you not go forward with?

33:24 CH: Yeah, so of the organizations that we evaluate, what percent do we not go forward with? It's a little bit of a challenging question 'cause I mentioned that we spend the most time with the organization that seemed most likely to succeed, and so it's hard to put a number on it. I think we talked to something like a dozen organizations this year. I might not have that number exactly right, but it is a pretty small number. We're not talking to hundreds every year. Through GiveWell's history, GiveWell has at least had preliminary conversations or looked at the website of thousands of organizations, but in terms of groups that we are actively in conversation with, as potential top charities, it's a pretty

small number because we're really just looking to move forward with those that seem most promising, and we also don't want to take up organizations' time.

34:13 CH: It's not a good use of our time or their time if they're unlikely to be a good fit. And so, it's a pretty small number that we have serious conversations with, and an even smaller number that we kind of get to interim review stage, which is sort of the next step of our becoming a top charity, and then a very small number that we get to full-on top charity review status.

34:33 CH: So it's not a very big number per year. And a lot of the winnowing of potential charities we could talk to happens as part of the fact that we do that program assessment first and say we're really looking for organizations that work on this set of programs. There are about 20 programs on the list, so it's not a super narrow set, but it's not a very broad set either. And so yeah, it's a reasonably small group that we end up spending time with. Great. And so it's a little after 7:50, so I think I should call the break here, but I'll be around for the rest of the night, so you can talk to me. My colleague Isabelle here can talk about GiveWell's research, and Caitlyn over there. So if you have GiveWell questions, feel free to come talk to any of the three of us during the break or after the fact. Yeah, we'll take about a 10-minute break here and then...