

GiveWell New York Research Event, June 11, 2015 –
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 Holden Karnofsky: Yeah, thanks for coming, everyone. I'm HK, I'm co-founder of GiveWell and I now lead the Open Philanthropy Project. So we basically have two very distinct projects going on, both of them related to figuring out how to do as much good as possible with your donations, and publishing what we find, but they have different emphases. And so, the first session is going to be on GiveWell top charities, the second session is going to be on the Open Philanthropy Project, and just so you know where I'm at, my focus right now is the Open Philanthropy Project. That's where I spend pretty much all my time. And for top charities, I am still in the loop, I'm still there, so hopefully I'll be able to answer questions, there may be some things I don't know. So, I'm going to start off talking about top charities, and that's what this whole session's going to be about. What the top charities are, for those of you who just want the basics, we try to find organizations you can give to that we consider proven, cost effective, and scalable, in the sense that they do something that we can find strong evidence, really helps people in a meaningful way.

01:02 HK: It's cost effective in the sense that you sort of get a lot for your money, every dollar you give goes as far as we can figure out how to make it go and scalable in the sense that they are going to do more with more money. So if we find a good organization that has plenty of money and doesn't have much plans for more, we're not going to recommend that. And sometimes we take an organization off of our recommendations because we feel they've raised as much as they're going to use productively in the next year. So we really do think on the margin, we think about what the next dollar's going to do and that's what we aim to do.

01:35 HK: The first thing I'm going to start off with is just our four current top charities, and naming these four charities is the biggest decision we make every year. We generally announce them in December, that's when people do most of their giving, and by now, we're at the point where we're moving tens of millions of dollars every year to these top charities. So it's very important to us and we not only are very careful thinking about who should be a top charity, but we also do a lot of updating. And every year we go back and we try to see how they're evolving, how they're spending their money, and really keep up with them.

02:09 HK: So the first thing I'm going to do is talk about our four top charities and what's happened, after a very basic summary, what's happened with them since last time we checked in, since we kind of made our announcement in December. And then I'm going to take questions on those, and then I'll talk about some of the other things GiveWell is doing, basically, to find future top charities to add to that list. So starting off with that, the top charities, I'm going to start with Against Malaria Foundation. That's the group that we moved the most money to last year, I had counted \$10 million total in 2014 that we could track going to Against Malaria Foundation as a direct result of our recommendations. As a quick refresher on what they do, they're a group that distributes insecticide-treated bed nets for preventing malaria. So basically, you live in a malaria-heavy zone with a lot of mosquitoes, you put a bed net with insecticide on it, it kills mosquitoes, it blocks mosquitoes, it reduces malaria. And we believe there's very strong evidence connecting these bed nets to lives saved and to improved health, and that they're very cheap for how much good they do and how much they help people.

03:16 HK: So Against Malaria Foundation is one of our top charities because the intervention is so good, because bed nets are so good, but also because we think they have really high standards around reporting and around sharing information about how distributions are going. So they not only will pay for the bed nets, but they will make sure that there is going to be a follow-up to see if people are using the bed nets, one month later, six months later, a year later. They will try to do what they can to get malaria prevalence data, which isn't always... You know, isn't always the most reliable data, but can give you some sense of whether the bed nets are working as intended. And compared to other organizations distributing bed nets, AMF, as we call them, is a good source of information and a group that really follows up what it's doing. So last year we moved about \$10 million to AMF and, a big factor last year, we've had some questions with AMF over the years about... It can be challenging for them to get money out the door because what they try to do is they try to come to an agreement with a government, generally a government somewhere in Africa, to participate in the national bed net distribution.

04:25 HK: And they ask for more than other organizations ask for in terms of wanting data, in terms of wanting follow-up. And that can make it challenge... It can make it a challenge to get a deal together. And so, we've had times in the past when we said, "AMF hasn't made a new bed net deal in a while, they're sitting on a lot of cash, they're going to be not on our top charities list for a while." And we have done that in the past. Last year, they got back on the list after signing a major deal in Democratic Republic of the Congo, and that's why we felt that we were ready to see more

money go there because we felt that they had spent a lot of the money they had. This year, we've just been keeping track of them, and so far, there's not a lot of news. We're still waiting on the first report from DRC, which is the country that they have the new distribution in. They haven't partnered on any new distributions, and so we're just watching and waiting. And if next giving season rolls around and there's still no new distributions and still no news out of DRC, we'll probably still recommend them, but we may limit how much money we're trying to move there, because we don't want a lot of money sitting in their bank account when they could be getting out into the field through some other vector.

05:37 HK: So that's AMF. GiveDirectly is the organization that gives direct cash transfers to very low income people in the developing world. So basically, they try to find the lowest income people they can, they give them a one-time transfer, generally in the range of a \$1,000. That can be very meaningful, that can help people sort of invest, gives them a lot of options for what to do. And this is under the theory that a lot of times, the people who know best how to spend the money are the people who need the help themselves. We've been recommending and following GiveDirectly since the day they were founded, and they've been growing a lot and we've been really encouraged by it. So last year, we moved them about \$9 million, they spent about 15 million, and, this year, they continue to expand their field capacity. They think they could spend another \$40 to \$50 million. So they have a ton of room for more funding and we've just been very happy with GiveDirectly at every step, just all the... All of the news coming out of them has been good, in terms of their ability to track who's getting the money, make sure it's the right people, the people who are eligible, the people with low incomes.

06:41 HK: Follow up, see if there are any problems, continually learn from what they're doing, be very transparent about reporting. To us, they're kind of a model charity. One other interesting fact about them, they've always been interested in doing this but I think recently, they've gotten more consistent about saying, "Every dollar we spend is going to be part of a research study." So for every dollar you give them, they're trying to get 90 cents directly into the hands of someone with very low income, which is great. Also, they're trying to make that be in the context of some kind of randomized study, so we can learn something about what the best ways to give out cash are. In other words, do you want to give it to a whole village at a time? Does that create village-level issues? Does that create inflation? Do you want to give it with a certain framing? Certain suggestions about how to spend it, just learning and constantly generating new information is, I think, a major strength of theirs.

07:34 HK: Next top charity, Schistosomiasis Control Initiative. They basically do mass administration of drugs that kill low-level parasites in children in Sub-Saharan Africa. So there's a very large number of people, something like half the population of Sub-Saharan Africa, at that age that has these parasitic infections. And it's not really clear what the infections do, they can be very... They can have no symptoms, but our read of the evidence says there's at least a pretty good bet that, actually, if you treat children and you kill the worms year after year, you can have a long-lasting impact on those people's lives, that they'll be earning more 10 years later. So SCI does that and last year, we moved them about \$6 million, which is in the range of what we thought they had room to spend.

08:21 HK: And this year, there's not a lot of updates yet, we've been following them, I think we've always had an issue with SCI, that we have more trouble understanding what they're doing than some of our other top charities. So it's just, "Where did the money go? What exactly is happening? How're things going?" has been tougher questions for us to get to the bottom of with them. It's been very consistent. We've had many different staff work with them over the years, so I think this is a fact about SCI at this point. And that continues to be an issue, so we don't have as good an understanding as we want, we're still working on it, they believe they have room for a lot more money and I think we're going to have to sort that out by the end of the year.

09:00 HK: Finally, Deworm the World Initiative. They also do deworming, same basic intervention, though they sometimes work in different areas, a lot of their work has been in India and last year, we didn't think they had a lot of room for more money. We actually think they're a great organization, we've had a great ability to communicate with them and get the information, and think they do well. But last year, we didn't think they needed a ton of money, we estimated they had a gap of about \$1.3 million over two years and we moved them about \$1.1 million just last year, so this year, it's been a little unclear. At one point, it looked like they could use a lot more money to go into Nepal, but then with the earthquake, that's become a lot less clear and so we're basically... I think their situation is that they are going to wherever they can find an opportunity to get a national deworming program up and running. And they see their role as advocating and providing assistance for the government to do that and they think that, in a lot of cases, once they get enough momentum in a country, another big funder will come in and help them but what we need to do is give them the basic funding to run their operations and look for those opportunities, and we've been doing that.

10:08 HK: So those are our top charities and I'm going to pause right there to take questions for a little bit before talking

about some potential future top charities. Yup?

10:22 S?: Have you been doing tracking of what impact GiveDirectly grants have on people's lives, and what that looks like on the ground?

10:29 HK: Sure. Impact of cash transfers, how does it affect people's lives? This is... Direct cash transfers is actually one of the interventions that has had the most study. There's been a lot of a studies on it and I think partly because it's controversial, I think a lot of people are not comfortable with this. A lot of people say, "Well, does it really make sense to give the money directly to the people? Can't we do better than that?" And I think the general picture that emerges from the study is a little bit ambiguous, but is largely positive. I think there has not... People have looked for and largely not found evidence of things like the money going to alcohol, the money causing problems, the money being misspent. People have seen significant increases in food consumption, food security, estimates of whether people invest the money and get a return on it have generally found that they do invest the money and they do actually get a really big return on it. And I think a lot of times, that's because you're talking about people who have so little and especially have so little cash and so little, just, liquidity, ability to take advantage of opportunities.

11:34 HK: They may be renting a motorcycle by the day and buying it or perhaps having to fix their roof every few months when it rains, 'cause the roof is made out of mud, buying a metal roof, can just, can produce some serious returns. So that's what the evidence has generally said about GiveDirectly, they're the only charity that we know of, at this time, although there's another candidate... Well, sorry, they're the only one of our top charities. A couple years ago, they were the only charity at all that had a randomized controlled trial of their own program. In other words, they're studying the very work they're carrying out, not just putting the old studies of people doing similar things. And that found, you know, it's not a very long term study yet, but it has found improvements in food security and improvements on a few other dimensions, and I think, for cash, it's particularly hard to study because it's hard to pick up an effect on any one indicator that's big enough to achieve statistical significance, because you're giving people money in a way that they can spend it however they want.

12:34 HK: So people might spend money in 10 different ways, get 10 different benefits, and you're going to have trouble measuring one to the point where you can statistically distinguish it from noise. But I would say, so far, the returns are positive. We do... You know, there's an argument that bed nets cost \$5 or so, including distribution. Deworming costs 50 cents a year or so, including distribution. You know, maybe, maybe you get more for your money with those than with cash transfers. Our best estimates to do that say, "Yes, you do." You might get something like five times more benefit per dollar with other interventions. On the other hand, our estimates are incredibly rough. And when you do an estimate that rough and it says '5X', I think you shouldn't be that confident, and I think there's a lot of other advantages to GiveDirectly, as well. Yep?

13:24 S?: Why are the direct cash transfers so much more effective than say, like, micro-finance or something, they've got the same thing?

13:31 HK: Sure, so direct cash transfers versus micro-finance. This is something we... We actually did write a blog post directly comparing these two things. You know, I would say, there's... I mean, there's a few issues. I think for one thing, micro-finance is generally considered inappropriate for people who are as low income as the GiveDirectly recipients. So, the people doing micro-finance, it's often the case that they're, you know, they're very poor by any kind of developed world standards but they're... They have kind of enough going that they can, that they can take a loan, they can pay it back with significant interest. And a thing to know about micro-finance is... You know, micro-finance being like small loans to people in the developing world, standard practice in micro-finance is high interest rates. The interest rate can be... We've seen interest rates of 100% a year and rarely seen interest rates below 30% a year. So this is like comparable to credit card debt and you need to make a serious return just to break even on these loans. Like I said, people can make serious returns on cash and, you know, a lot of times people have so little cash that they...

14:36 HK: A little bit goes a long way and they can make a serious return but in micro-finance, you need a serious return just to break even, and I think break even is a lot of what we've seen from micro-finance. There's been a lot of studies on micro-finance, as well; unlike the cash studies, they generally come up something approaching empty-handed. They haven't really found harm, they haven't really found benefit, it's been a very muddy picture, and you can't really look at the studies and say that this helps or hurts. Remember, I mean, loans could hurt. They have high interest rates and certainly high interest loans hurt a lot of people. I think when you see a lack of evidence, I think that kind of is consistent with my picture, which is, on average, probably some people get help, some people get hurt and, overall, on average, you probably see something like break even. And it's cool that people are building institutions that provide

loans and can support themselves, and kind of become sustainable, but as charity, yeah, I would definitely go with the cash for all the reasons I just said. Yeah, Justin?

15:32 S?: In selecting the top charities, how do you weigh the benefit... That immediate benefit versus choosing charities that are modeled in terms of transparency and sharing their research?

15:43 HK: Sure, so in choosing top charities, how do we weigh the immediate benefit versus the benefits of holding a charity up as a model? So, I mean, we care a lot about... I think we really put the direct benefit first, that's the main thing we're looking for. I would say that if we... I would say that in cases where there's an extreme difference, you'll see us put some weight somewhere. So, in other words, our top charities are generally, I would say these are as cost effective as anything we can find. Period. Like just any charity we're aware of, in terms of how much good you're doing for your dollar. Now I can name you a lot of other charities that do equally good interventions. So, for example, and I will talk about this a little bit, but immunizations to fight maternal and neonatal tetanus, there's a group that does that. There's groups that do nutrition. Like, there's interventions that are as good as our top charities. And generally our rule is, if we have some charities that do the good interventions and we've investigated them, we know everything about them, we've gotten all the information from them and we're able to track how things go and learn from them. And we've got some other ones doing equally good with interventions we know kind of nothing about them, we're going to recommend the first crew. We do publish a list of the second category for people who want that and that is on our website. If anyone can't find it, you should email me.

17:03 HK: But that's generally how we do that. And then, if we saw something that we thought, "Okay, this could be 10 times better than our best charities now," I think, there, we would start to relax the requirements for transparency and be a model and things like that. So, for us, you know, transparency and being a model is a little bit more of a tie-breaker. Though I also think that in the case of GiveDirectly, they are a noticeably better organization than... Like in terms of just like, how consistently they track and report what they're doing, how well they've grown, and how much we feel we really understand what's going on there, I think they're noticeably sort of a standout there, and I think that is part of the reason that we keep them on the list, even with our calculations saying they may not be quite up there on cost effectiveness. Yep?

17:51 S?: How do you think about, I guess, the time dimension of cost effectiveness, when you try to compare it to... Sort of for different causes even before you get to a specific charity...

18:01 HK: The what of cost effectiveness? Sorry.

18:02 S?: The time dimensions.

18:03 HK: Okay, time dimensions, yeah.

18:05 S?: So say there's some structural investment you can make that won't pay off for 30 or 50 years or something like that, but you still very high confidence it will, it seems like there's sort of a very, there's a high level discretion and subjectivity to deciding how to think about that today, versus, "I'm going to save someone's life" statistically over the next two or...

18:26 HK: Oh yeah. I mean, you run into situations where you could, not hypothetically, but really, you could give out bed nets and save lives in the next few years, because those are primarily infants. Or you could give out deworming and figure that most of the benefits are maybe 10 years from now, when people are out of school. Or you could give transfers that are cash, that's kind of a mix that people can invest and consume. And one of the big debates in these circles is around, "What discount rate do you use?" The discount rate being, "How much more should we value benefit today than benefit 10 years from now?" All I can really say is that when we run our calculations, we use multiple different discount rates, and this is a common thing about our calculations, is we do them lots and lots of different ways.

19:14 HK: So we have a spreadsheet, and it says... What we actually do is, we have a spreadsheet that says, "Here are all these variables and we don't really know the right value. Here's a few values that we think are reasonable, that reasonable people could defend." And then Elie will fill in what he thinks his best guess is, and I'll fill in mine, and Alexander will fill in his. And then we'll get a whole bunch of people and we actually publish a spreadsheet that's like, "Here's HK's, here's Natalie's, here's Elie's." And a lot of times... A lot of times what happens is like, Elie and I will disagree here in one direction and there in another direction, and our ultimate estimate won't be that far off. So our answer to these very thorny dilemmas is not to try to debate them to death, but rather to try and put the information out

there, try a few different things and have conclusions and recommendations that are kind of robust, that don't... We don't want to make recommendations where it's like, "This is the right thing to do. Oh, unless you don't have a 7.14% discount rate on cash transfers or something." So yeah...

20:14 S?: This is well, a follow up, I guess, I've been a little surprised... Not surprised, I guess, but noticed that it does seem like you guys go towards more immediate things. You know, 10 years down the road, than say, structural things that, building some infrastructure that maybe has a 50-year payback period, something like that. Is that... I guess across the range of parameters of variables you used, do you find that you sort of still, in order of magnitude, that included the immediate stuff, is that kind of implicit in what you're saying?

20:41 HK: Sure. So why do we do immediate stuff instead of long term stuff, like building infrastructure? The way I put this is, our Top Charities Project is all about finding the organizations we can be highly confident that they're providing excellent value for money. And highly confident is important. So we look for things that have evidence bases, we look... Bed nets, there have been over 20 studies in which they randomized who got bed nets, they looked at what happened. We can pool all the studies and we can understand that bed nets help people and we can believe that they'll help people again when we give them out again. So that is how the Top Charities Project works. Now, you might say, "That might not be the best way to give. Sometimes you don't know the answer, you have to take a risk, you have to take a guess, but you do a lot more good." I would agree with that, I would say that that attitude brings a lot of new options and it brings a lot of new challenges and a lot of new problems. That is the Open Philanthropy Project, essentially. And so that's going to be the topic of the next session.

21:36 HK: And I think as people will hear, once you stop looking for something reliable and you just try to do the best you can, you end up in a world of massively more judgment calls, massively more guesswork. I'm going to say thing after thing that people are going to say, "Why is he saying that? That sounds crazy." So there's disadvantages, but we want to do the most good we can. We think both approaches have a lot of merit, and we do both, but they are different. And so that's the subject of the next session. Yep?

22:05 S?: What's the... So, in terms of AMF, my understanding of the Congo is that it's a place where the rule of law is not as strong as in other parts of the world.

22:14 HK: Yup.

22:15 S?: What are the risk of, like... What's the worst case scenario of that distribution?

22:21 HK: Yeah, worst case scenario working in DRC. It's a major concern of ours. I mean, I think this is a huge challenge and, if we can get it to work, it would be amazing to have a functional bed net distribution that we're able to learn from and get the data from in DRC. If it doesn't work, what's the downside? I mean, I think the worst possible case is that the money is basically lost, that the bed nets don't get distributed. I think that's somewhat unlikely, I think that a decent number of those bed nets will get distributed. Now, will all of them get distributed? Will a lot of the money be lost to kind of who knows where? Will we ever get the data on how the bed nets were used? That's exactly what we're waiting to find out, and we followed the negotiations with DRC and the agreement that was signed there, and we looked at it and we said, "This is something that might work. Let's see how it goes." But yeah, I mean, it's a question. And it's a question we've got our eye on.

23:15 S?: So the government has to send people into the villages, right? After every three months or six months? And so they can just decide at any point, like, "This is a pain in the butt, we don't feel like doing this anymore."

23:26 HK: Sure, I mean, it's a negotiation. So the question is kind of, "Does the government decide how they're going to collect data and how they're going to report?" It's a negotiation. So AMF says, "This is what we want, if we provide a large amount of monetary value in the form of nets," and the government says, "We can do this," or, "We can't do this." And then, if they eventually don't follow through on what they said they would do, I mean, I think AMF says, "We're not going to work with you again." And I think, AMF, to me, is fighting harder for the things that... For the follow-up, for the data, than other groups distributing nets. But yeah, I mean, there's no guarantees here. And I should say, "Look for the most reliable charities we can find," but there's a ton of risk here, and I think there's all kinds of uncertainty in these top charities, and I think they are the best we can do on that front, but there's still a lot of unknowns.

24:14 S?: One thing to add, I just want... I know that AMF negotiates with both the government in doing the program, the National Malaria Control Programme for each country, but also we've looked at independent organization which

usually is the implementing program. It's not actually government, like reports who might, who would go look at the data. I'm trying to understand their model [24:29] ____ early. I don't know who the program is with DRC but it's going to be the organization that's used to work in there, I think that's [24:35] ____.

24:36 HK: Right, yeah. I mean, it may be instead, of the government directly, it may be the gov... It just has to be that the government allows the non-profit to go in. In some cases, allows a non-profit to retain government employees for this purpose. There's a variety of ways of setting it up and I had forgotten that point. I know... I know that's how they did it in Malawi, and I actually don't know for sure that's what they're doing in DRC. Yeah, Jason?

[background conversation]

25:16 HK: How does AMF fit in to global bed net distributions, are they just a drop in the bucket? There's certainly a lot being spent on bed nets and a lot of it's by government, so it's in the hundreds of millions of dollars a year, but there's still a massive gap to cover. We estimate this every year, we say, "How much would it cost to cover everyone who theoretically needs a bed net this year based on when the last time they got a distribution was and based on how many people are at risk of malaria," and it costs a lot. There's a gap of hundreds of millions of dollars as well, so there's hundreds of millions going in, hundreds of millions gap. And AMF at this point is able to come in with something in the range of \$10 million, so they're not enormous, but they're not a drop in the bucket. When AMF was smaller, they would work with a non-profit, they'd go to some little districts, they would do a little distribution, it wouldn't necessarily even be coordinated with the government. I mean, now they're at the table and I think they have real leverage, but they want more. I mean, AMF says...

26:13 HK: I mean, I think AMF's opinion is that we should recommend them enough money that they could cover the entire country and that would greatly increase their negotiating power. And I think it would, but our position has been... We want to see consistent performance at the level they're at, before we recommend that kind of money. I'm going to go to the next topic, when I break for questions, people should feel free to ask about anything, but I want to make sure I talk about what we're doing to find more top charities. So this is something... We have four standout charities that we announced last year and, in some of those cases, we think those are groups that might become top charities in the future. So one of them is Development Media International, and they basically put out mass communications in the form of things like soap operas or radio ads that are trying to promote healthy behavior. So, for example, how to know when your child needs to go to the clinic or the hospital, and making sure you do that and things like washing your hands. I mean, these are simple behaviors that can save lives that you can transmit to a really large number of people per dollar when you use mass media.

27:15 HK: It's an interesting model and one of the things we like about them is they have an ongoing randomized controlled study of their own work. Last year, the preliminary results came out and it looked kind of preliminarily promising, but, like, it was just... All it was, was data from surveying people about how they'd change their behaviors, and that's an obvious problem if you're measuring the impact of kind of telling people what behaviors are desirable. This year is when we're due to see the actual mortality data. So they'd been somewhere for two years, giving their message, what happened to child mortality? We're expecting those results this fall and as soon as we get them, we'll have a lot more opinions about how much we want to recommend DMI. Another group, so two other groups that are standouts, GAIN and IGN are both nutrition groups so they work on fortifying food with important nutrients that can make people healthier, and we've done research on those and think they're good interventions. And these two groups are major priorities for us, I think they're very hard to understand because they are in the technical assistance advocacy model.

28:19 HK: So there are more groups that are trying to help governments implement programs, and that is harder for us to evaluate than groups that just kind of fund things directly. So we need to understand case studies of where they've had an impact and we need to ask people, "Did they really play this key role?" That can be a very hard thing to do, but it's a high priority for us and they may be on the recommended list at the end of the year, they certainly both have a lot of room for more funding, and they could definitely use donations. And then, finally, Living Goods is a standout but we have not prioritized them. They're a group that basically markets potentially life-saving goods and we estimated that their cost effectiveness is good, but it's not as good as our top charities and we don't think they're comparable to GiveDirectly on other dimensions, so they're probably going to remain a standout. Then there's some groups that are not standouts that we think have potential. So UNICEF runs a program on Maternal & Neonatal Tetanus, and they basically immunize every woman of child bearing age with three rounds of tetanus immunization to prevent the transmission of tetanus and a death in childbirth or just a death early in childhood.

29:31 HK: And we think... UNICEF is a tough group to work with, it can be hard to tell if they really need more money because they're a massive organization, and we don't know how the different parts of UNICEF relate to each other, but we believe that they really do... We've been hearing the same story from them for years now, that they wish they could go into countries A and B and they still haven't done it. So we think there's a lot of promise here, we think it's a good intervention. We think that the intervention depends dramatically on exactly where they're going. So, this is... This immunization stuff can be way better than bed nets or way worse, and it all comes down to where they're doing it and what the existing levels of tetanus are. So we really need that information and we are trying to get it from them, and I think it can be hard to work with UNICEF, but we're... It's a high priority for us and we're trying to get there. They may be recommended this year. We've looked hard for more deworming charities 'cause we love deworming but we feel like SCI and Deworm the World both have had limited ability to get a lot of money out the door to cover the whole world, and we've had our trouble communicating with SCI. So we have looked hard for other groups that do deworming.

30:40 HK: It hasn't really gone to there yet, but we've kind of looked at all of them and we're still thinking about which ones might have potential. We've also looked at other groups that work on nutrition, work on bed nets. We really tried with the bed net organizations to just say, "At this point, we can move a really large amount of money if we can just see more accountability in terms of who's using the nets, how it's going." So far, not a lot has come of it but we're continuing to work on it. Then there's even more further out work to try and see more top charities in the future. So we have been trying to think. We've tried getting money to the best, we've tried investigating the ones that are kind of almost there or might be there soon. What can we do to create more top charities or help more top charities come into being? We have provided early funding to a group called "New Incentives" that does conditional cash transfer. So that's when, for example, instead of just giving someone cash, you give them... You pay them for every time that their child demonstrates enough attendance in school or every time that they are demonstrating their child has had their shots correctly. And so you're trying to get a double benefit there.

31:47 HK: A lot of studies of conditional cash transfers, looks good. This organization, I think to become a top charity, needs to build a really robust operation where they're working in difficult places and executing this program reliably. And so we have provided them money to try and do that and try and demonstrate they can do that, and we gave them a grant a couple years ago that we renewed this year, but it could be a while before they're there. Another thing we've done is we've tried to think about, "Okay, it's not that common to see a charity that's really excited to evaluate itself, but what if we paid for someone else to evaluate it?" So we're starting to experiment with that, we're working with a group... SCI is the group that is our top charity that we have been sort of least happy with their self-evaluation and we are now working with a group called "IDinsight" to experiment where IDinsight would go into a country and provide monitoring data that we would find more informative than SCI. So immediately after deworming, asking people, "Did you get the pill? What did the pill taste like?" And things like that, to make sure this deworming came as intended. And if that worked out well, that could go anywhere.

32:55 HK: You could imagine taking any group that does a good intervention and saying, "We want IDinsight or someone like that to do the evaluation." And then there's stuff that's even more kind of blue sky, so there's... What are the... So far, I've talked about charities that do good interventions, but we wish we had more info about them. What about interventions that could be good, we just haven't studied them enough yet, we don't have enough research. So we've been experimenting with Evidence Action, which is the same group that runs Deworm the World, on this and so there's basically... We've experimented with a few things but this is the main thing in this category that we are excited about right now. There's one study implying that programs focused on helping people or encouraging people to migrate from rural to urban areas can result in people making a lot more money, and being a lot better off and actually, though they're... The first time, they were kind of nudged or incentivized to migrate. The next year, they go back of their own volition... I mean, their own volition either way, but they go back without the same nudge or subsidy.

33:57 HK: So it's an interesting study and then there's a question of, "Is this the intervention that only worked in one place one time or can this work at scale?" Evidence Action is trying it in a larger scale in Bangladesh, 4000 people, compared to something like 400 for the study. They looked into doing this in Africa, concluded that the place they looked, I think it might have been Zambia, but I don't know for sure, they concluded that it wasn't going to work there. It just wasn't the right dynamic in terms of the rural and urban areas and the job opportunities. And then, if it goes well in Bangladesh, they're going to try and scale it really significantly and that's funding that we might recommend. That's funding they may also just pull in from somewhere else. I mean, there's other funders and if they get people excited about seasonal migration, they may just... We may end up saying, "Well, we failed to create a top charity because they got funded by someone else. There's no room for more funding," but I don't know. That still feels like a good thing to have done if it works out. So that's our kind of experimental work.

34:53 HK: And then the final thing I want to say is just how much our capacity is increasing? So for many years, I've been kind of showing up and getting questions like, "Why haven't you looked into this? Why haven't you looked into that?" And I said, "Because we don't know enough people, we don't have enough time," and that's always going to be the case. We're always going to be limited in what we can do but it is changing in a positive direction. We've been hiring a ton of people, we're about 30 right now, and we've also just been building the organization into something more robust, less dependent on me and Elie. So actually, I am barely involved in the top charities work anymore. I focus on Open Philanthropy Project, which is the subject of the next session. I'm obviously... I'm engaged, I'm talking to people about it, I'm giving opinions on it, but it's not something that's really depending on me. And we think that's a great thing and, you know, now, we're taking... We're basically trying to build an operation that can run itself more, and can do more than Elie and I ever did when it was like just me and Elie, and do it without even needing that much involvement from us.

35:54 HK: So as an example, Natalie Stone Crispin has been with us since 2009. She now manages several people and has been taking just a lot more responsibility in all this than we've ever had someone else take before. She's coming back from maternity leave soon, so we've been in kind of a transition, but that's an example of something that I think has been going really well. And sorry, when I said we're 30, I meant, by the end of the summer, we expect to be 30. We have a lot of new people coming on who have accepted offers. So, I mean, when I list all of the things that we're doing and I say, "We're looking at this charity, we're looking at that one, we're doing this experimental thing, that experimental thing," that is something that we could not have done a few years ago, even though a lot of our most senior staff from a few years ago are now not even available for this work anymore. They're working on the Open Philanthropy Project, so it's been... You know, it's been the ability to have a whole new project, which is going to have its own session, you'll hear all about it, but also get better instead of worse, at doing a lot of stuff on top charities, that's what we get for hiring all these people.

36:54 HK: And as you grow, a lot of issues come up, so just building the organization as an organization has been a major focus for us, and something we've put a lot of time into, so that's also something I want people to be aware of. So yeah, I'll pause again and take questions for the next 10 or 15 minutes, and then we'll break. Yep.

37:12 S?: Thanks a lot, HK. Things have been really interesting. I've got two questions, the first is, whether you guys have been following the ultra poor graduation model? It's a big study. For people who aren't familiar, that just came out recently. Six different RCTs with the same program around the world with seemed to be pretty encouraging results. I would be interested to hear if you guys are following that, thinking about the cost effectiveness of that [37:36] ____ cash transfers on the radar. My second question is, now that you're branching out a little bit outside of the public health world and your traditional recommendations, if you're also thinking, again, about education interventions? There have always been a few bumps in RCTs in this space with [37:57] ____ remedial education and some other programs, so if that's going to be back on the radar at all?

38:00 HK: Sure. So I'll take them one at a time. Question one... Actually, I'll take question two first because what I actually did just now is I asked the San Francisco office where we are with graduating the ultra poor, because I think I... I have some idea of where we are, but someone else will know more, so this is an example of how things are not only not as reliant on me anymore, but I often don't know what's going on.

[laughter]

38:24 HK: So we'll see if anyone writes back, but anyway, on education, yeah, we haven't gotten there yet. I mean, we're still trying... We still have a basic intuition that this health stuff is probably the best bang for your buck, you just look at what it is. It's like a 50 cent pill, it's been studied, we know it kills this infection that a huge number of people have. We think it's going to generalize, it's incredibly cheap, you know, same with the bed nets. I think cash transfers are kind of a special intervention, but have a lot of their own appeal. So our priority has definitely been health rather than education. Of course, deworming, and to a lesser extent, bed nets, are believed... And to a lesser extent, cash, I mean, all these things affect people's ability to get an education. There was... One of the major research groups tried to compare a lot of different ways of getting kids to attend school more, and they concluded that deworming was the best or second best way to do it. And so it might be one of the best educational interventions in terms of what you get for your money.

39:23 HK: So a lot of things do work through that vector, and I would expect... I think there has been a lot. Last time we looked at education, the research really wasn't there, and I would expect, there'd be more research now, but I would also expect that we're still not going to find interventions that have really robust evidence base where we really think we can

generalize that it'll work again and it's really good cost effectiveness that's up there with like a \$5 bed net that protects you from malaria. So we may get there at some point, but we haven't gotten there yet.

39:53 HK: On the Ultra Poor Program. So the Ultra Poor Program is cash plus various other things like technical assistance and just helping people, trying to find other ways to help people make the most of the cash they get, and it's kind of a hybrid intervention that there've been a lot of good studies of, and that's something where I said I don't know exactly where we are.

40:16 HK: Let's see, I mean I think what we've been doing is just trying to get on top of all the research and understand what it says. So, basically, this is kind of a... It's a combo intervention, it's kind of a tough intervention to describe, there have been a lot of studies of it recently and a few years ago, the studies weren't done and now the studies are wrapping up. So I think where we are is, we're trying to understand the evidence and then if we like the intervention, then we would look harder for charities that work on it. Yep.

40:45 S?: Maybe you can comment on the idea of overlap, like you can incentivize, take your worm pill when you distribute a bed net, a lot of this seemed to... You're in position 'cause you have this overview of all these charities, where you can cost effectively make use of where they might overlap.

41:10 HK: Sure. Do we see opportunities to make use of overlap? We actually wrote a blog post in something like 2008, which is just ages ago, maybe 2007, where we said, "Hey, why don't people give out the bed nets and the pillows and the vaccines at the same time?" That isn't something that we've really focused on since then. I think for the most part, it is harder to combine these things than it probably sounds, so the bed nets kind of tend to come every few years and you try to make sure everyone in the community gets a bed net and knows how to use it. The deworming, you want to do every year or sometimes every six months, and that's really targeting school children. It's children of a certain age, you want to make sure they take the pill. Vaccines are generally for much younger people and so, yeah, I haven't seen... I will say that we haven't seen a lot of situations where we've been really convinced, "Oh yeah, we should put these two interventions together." But if we did see it, I mean, I hope we would remember to point it out because I think there are cases of this, and I think Nothing But Nets, at one point, was trying to piggyback nets on vaccine campaigns. So it's not...

42:12 HK: Or we've seen cases where deworming is piggybacked onto a general child health day, where people get a lot of different services. So it's definitely something that happens, it's not something where we've really said we see something that ought to be happening in that category. Yeah, Tom?

42:27 S?: I guess sort of a related point, years ago, VillageReach was a top charity and that was sort of a supply chain solution for delivering emergency care to people. So that's sort of a similar thing. Is that something you guys have investigated at all? Was there any sort of learnings from the experience of working with VillageReach?

42:46 HK: Sure. Maybe I'll just give a little VillageReach update. I mean there's not much of an update, Village Reach was the first charity that we ever got significant money to, and they try to improve logistics systems, so people can get immunized better. And we felt that we got them enough money to go forward with their program. And then, at one point, things were kind of getting slow and the program wasn't rolling out and they didn't really... You know, they didn't really make a push to us for the case for more money and so we parted ways at that point. We actually haven't been in touch since then. I did see that they got a grant from Gates Foundation that, to my eyes, looked like something kind of similar to what we had supported them to do. So, I don't... I think what they do is actually kind of a significant interest to some other funders, which is improving logistics for vaccination. So I actually wouldn't be surprised if they're doing well, but they're not a group that we've continued to interact with, in terms of like them coming to us for help with the funding gap, and so that's pretty much where that is.

43:43 HK: I mean, yeah, I mean, I think if we... If we saw another group like that, we would be potentially interested. I think the big hurdle is most... We see a lot of groups that have like an interesting model for improving health that's a little unusual or a little unconventional, and most of those groups just don't have data that's going to get anywhere. A lot of them say they do, and then we look at it and they don't. We've had a lot of that experience. Village Reach really did, I mean, I think they really had the goods and their data was something that convinced us, and I look back at it, I still feel that way, but it's not the usual situation. But I think we'd be open to it, we'd want to see not only the model but the data. Yeah, Ben in the back.

[background conversation]

44:49 HK: Sure. So let's talk about giving to GiveWell itself. That's basically the question, what's the story update on that? We generally put out a public update every year on where GiveWell is at, in terms of what our projected revenues and expenses look like and how much money we need. For the last few years, a few things have generally been true about our fundraising. Number one, we don't put it front and center because we don't need to and because we think it helps our credibility and it helps donor comfort that we don't go out and say, "Give to us." We go out and say give to our top charities and it's only once we... Generally, it's only once we've built a relationship with someone, we've talked to them a couple of times and we feel like they'll be interested and excited to hear this, that we talk about giving to us. Now we do make the blog post and we say, "We have this funding gap," we don't want that to be our top message and so that's generally not our top recommendation. So that's been generally true for a while. We have talked about putting something in... It's even possible we have and I don't know it, which would be funny, but we have talked about putting something in where, when you donate to a top charity through our website, we ask you to consider 10 percent, and that would be a box you have to proactively check.

46:00 S?: We have that.

46:01 HK: And so we have that, excellent. Okay, we finally did that, right. So I think that's a fair thing to be doing, like that's a common practice in organizations that try to help people give money, and I think that's something we ought to be doing and I guess are doing, awesome. To ask people, "Hey, you're giving to charity 10 percent to the people who helped you decide. Consider it." It's not like something we require at all or nudge very hard. So that's been true for a while. Another thing that's been true for a while is that we... Our general MO is to like... Our general MO is to try to not have a projected funding gap, so we try to get our revenues in line with our expenses. And usually, once we feel like we're there, we will be less... Just put less effort into trying to raise money for ourselves. So generally, what's happening is we look at where we are and if we see a big gap, we say, we have to call more donors and ask them to consider giving to GiveWell. And we're going to start, always, with the people who we feel know us best and who will be most receptive to that message, and then we move down the line. And then when we feel like we're fine, we'd just spend less of our time on that.

47:06 HK: And so what ends up happening, if you give to GiveWell... Like what do you get if you give to GiveWell unrestricted? I would say that what you're accomplishing is you're funding our operations, you're largely saving us time, you're often fungible with other donors. So a lot of times, it's like because you gave the X dollars, there's another person that we never ask for money and they're giving X dollars to top charities. So I think that's actually a good thing, it's a good deal in that respect, in that, when you give to us, it just means... It's like giving to our top charities in a way but it does save our time. It means we close that funding gap faster. And a final thing to note is that we have a policy in place to ensure that we don't just start piling up massive cash reserves. And so, if we ever got so much money that it was just way more than we needed by any reasonable accounting, and we have a formula for this, then we would be required to grant some out to top charities.

48:00 HK: We also have some other measures. I think right now, we've recently been deciding, or decided, to not take more than 10 percent of our budget from one donor. So if someone gives us more than that, the rest of it gets granted out to top charities. And that's for our robustness, so we're not too reliant on one source. So you know... But the... The thing about reserves does not tend to get triggered, like we don't go around raising money 'til we have a massive reserve pile, we raise money until we don't project a deficit anymore. I don't know that that policy's ever been triggered, but the fact that it's there is some comfort, that we won't just pile up money forever. So, you know, the bottom line... I mean, I think giving to GiveWell is a great thing to do. I mean, I personally do it, but it's not something... It's something that we encourage people to think about if they want to think about it. It's also not something that we try to aggressively push on people, and we're really happy when people give to our top charities. That's the front and center ask.

48:56 S?: Just a little bit more about the people, so you're expanding it to 30 people that come. Is it easy to find great people? Like, how do you...

49:04 HK: Yeah. Talk more about hiring?

49:07 S?: Right.

49:08 HK: Sorry, I want to say one more thing on fundraising, which is that, right now, our expenses are something like evenly split-ish between GiveWell and the Open Philanthropy Project, and that is something... You know, we're actually hoping to turn those into two different organizations at some point, but the way things work right now is that

the expenses are fungible between the two, so when you give, you're supporting both. And one of the advantages of that is that the Open Philanthropy Project has been... It's a young organization, or it's a young project, and so a lot of times they'll say, "I really need help with X. Elie, is there anyone who can do X?" And he'll say, "Yes, this person can work on it." So we have the staff going back and forth between the two, and Elie and I make those decisions about how to allocate people. That's important to know if you're donating. I mean, I think if you, at this point in time with how we run things, if you don't like Open Philanthropy Project and you don't want to be a part of it, you could earmark your donation for GiveWell, but first and foremost, you just need to be aware that this is an issue and you should let us know 'cause right now that is how we're running things and so we get the support for both from that.

50:12 HK: So Jason's question about hiring, how has hiring been? How has it been for us finding people? This is something I wish we had the time to write about, and I'm sure we eventually will, 'cause I think we've just learned a ton. A really brief story that might just help people hear a little bit about what it's like. I mean, when we started... I think we made our first hire in 2007, and it was me and Elie and one other person, and we just... We were not ready for a hire, we didn't know how to manage that person, we didn't know how to evaluate their work. We would say, "Hey! Go look into this thing," and then they would kind of email us something and they'd say, "I don't really know what to do with this." One of the things we've learned over time is that it's really important when you hire someone to have a good plan for what you're going to ask them to do, and how you're going to know if they did it well, and what their path looks like from there. The ideal is that someone comes on, you give them something to do, they do it, you look at it, you definitely know if they did it well, you definitely know how relevant it is, and then you can keep challenging them more and more.

51:13 HK: And then each step of the way, you retain your ability to evaluate them and give them a good experience. And that is a hard thing to do with GiveWell. It is, I think, a special challenge that not every organization faces, because if you're running a company where what you need is, for example, software engineering, that's a thing. You can look at someone's resume and see how good a software engineer they are. You can look at their work, you can see how well they did it... You can see how well the thing runs. With GiveWell, this is a much more ill-defined process and we've kind of made it up in many ways. Just how do you look into a charity, and what counts as a good job, looking into a charity? So I think that's something we've really had to work on and really had to develop, and I think it took us a long time to get there. And, you know, when we were really improvising, and it was me and Elie just kind of using our intuition, and saying, "Here's my next question, and here's what this paper said."

52:06 HK: We were not ready to do that. And what we've done is we've worked on that enough to understand what that process is like and to know what the steps are, and what good work looks like, and how to evaluate work, and I think we've really reached the point now where we really do have that down. So I think now we're able to hire people and we know... You know, we have a good path for them, and we have a good place for them to come in. And we also have a lot of other things going on to make sure they get oriented right, and that they come in, they know what's going on, they get to meet everyone. So there's a lot of stuff you have to do, I think, to make that work, and I think it's been a long process and one of the major things we've accomplished, to turn GiveWell from something that was really me and Elie just going... Doing what seemed smart to us, to an operation that's really able to have... To run in a more systematic way. So I think that's one of the things we've worked hardest at, and made the most progress on, and it certainly doesn't happen by itself.

[pause]

53:00 HK: Yep?

53:01 S?: Along the lines of you guys trying to go around with more people, and your dissatisfaction with a lot of the deworming initiatives, have you guys thought about starting your own little section that does that, and uses all the best practices that you guys have found, rather than trying to find a company that does exactly all the things?

53:22 HK: Sure. Have we thought about starting our own charity, perhaps a deworming group that works the way we want? No, I can't really say we've thought about that, that's not something we're planning to do. I think we don't... I think that's not our core competency, it's not really our skill. I think we have built up a lot of, I think, know how about how to evaluate organizations but that is not the same as how to run an organization. So I think we'd be... There'd be a very steep learning curve for us there, and I think it'd be a very different lifestyle as well. And we have a lot of ambitions and we have a lot of priorities, and so I think that one's far enough down the list that we're not going to do it. We'd love to see someone else do it and in some cases, we have. I mean, I think GiveDirectly, you could characterize that way and I think New Incentives, you can characterize that way. So I think I want to take maybe one more question,

if there's one more, and then we're going to take a break before the Open Philanthropy Session and people are welcome to leave then if they want. Yep?

54:18 S?: So is it your ambition to get larger and larger, and allocate a larger and larger share of the world's giving?

54:25 HK: Is it our ambition to get larger and larger, and allocate a larger share of the world's giving? I would say, mostly yes. Certainly, we want to move more and more money, and the bigger that number gets, the more it becomes the case that we can call up UNICEF and say, "This isn't going to be a drop in your bucket. This is major money. This is worth your time to really engage with." So yeah, I mean, we are trying to move as much money as possible to the best opportunities possible. I think there is such a thing as just growing your staff beyond where it needs to be. I mean, we pay attention, how much do we spend on our operations and how does that compare to how much money goes to our top charities? And if that number got unreasonable to us... I mean, right now, it tends to be around \$10 for our top charities for every \$1 we spend on ourselves. If that number looked more like two to one, I think we'd probably be spending... I mean that might be fine in a short term as we try to grow but I think as a long term ratio, that's not... That wouldn't be good.

55:21 HK: We could probably trim our staff and spend a lot less money, do almost as much good. So I don't think it's... We think there's such a thing as too much of that but, right now, we're in growth mode for sure and especially because Open Philanthropy is in major growth mode. Open Philanthropy feels like we cannot find enough places to put all the money that we could move if we could find the right place. We need way more people, we need way more knowledge, we need way more opportunities. And so, we keep taking people at GiveWell and asking them to work on Open Philanthropy, and then that keeps creating a challenge for GiveWell to just stay in place by adding in new people who can replace the old ones. So, for now, we are aggressively expanding, but that's not like an iron law. That's just what we're doing 'cause we think it's the right thing to do.

56:07 HK: Cool. Well, thanks very much everyone and let's take a little break.