

**GiveWell San Francisco Research Event, November 18, 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.

0:00:00 Elie Hassenfeld: All right, well, thanks everyone for coming. I'm Elie Hassenfeld. I'm one of GiveWell's co-founders, and this is...

0:00:08 Holden Karnofsky: Holden.

0:00:09 EH: The other co-founder. And so we really appreciate you coming out so we can have that research discussion tonight. Just a couple preliminary announcements. We're recording this. As you probably know, we are a transparent organization. We try to make as much of our research public. And so we're planning to post a recording and a transcript of this event on our website. If you speak up and ask a question and would rather we not include that and what goes on the website, just let someone from GiveWell's staff know after the fact, and we're happy to pull it out of the recording or the transcript. So please don't hesitate to speak up. We're happy to edit things out if that's helpful. You may also notice that there's a couple photographers here tonight. They're just trying to get a picture of sorta the research event for the purpose of an article that's going to run in the Chronicle of Philanthropy. Primarily, they're trying to get a picture of Holden and me.

0:01:02 EH: They might take pictures that includes some of the audience members. If anyone feels uncomfortable with that, feel free to let us know after the fact, or I guess, move more in the direction of the back.

[background conversation]

0:01:15 EH: Cool, great. So, we are... The basic way that this is going to work maybe familiar to those of you who've come to one of these before, we're going to go through and in the first hour talk about GiveWell's top charities and the research we've been doing there. In the second hour, talk about the Open Philanthropy Project and research that we've been doing there. Throughout the hour, we're going to give brief updates on various things we've been doing. And we'll pause and take questions. And the goal of having you ask questions is not for people to ask things that we've never said before or you couldn't possibly find on our website. We know that our written materials can be somewhat hard to parse, and so we'd really encourage you to ask whatever questions you have on your own mind, and give us a chance to try and help you understand all the work that we're doing. So, the sort of big news of today is that we're very close to releasing our updated charity recommendations for GiveWell. We at GiveWell tend to work on an annual cycle, where we're trying to refresh our recommendations so they're most up-to-date near the end of the year since so many individual donors give significantly during the holiday season.

0:02:33 EH: And so we are planning to release these updated recommendations within the next few days. There's not a lot of news on that front in terms of the organizations themselves. We actually anticipate having the same four top charities and the same four standout charities this year as last year. But there are some relatively major updates in the overall situation of each individual charity, and our guess about their future prospects, and also the research that we've done, and so we're excited to talk that all through. I hope that you'll resist the urge to live tweet this event so we still have the opportunity to post the news on our website in the next few days. So without further ado, I want to jump in to the substantive updates. So I think there's a couple big pieces of news about our top charities, and I think that the biggest one is that, there is significantly more room for more funding among our top charities this year than there were in past years. And so room for more funding is this concept that we use try and determine how much, how well an organization can use additional funds that it receives. Will it be able to put those funds to good use and have high impact.

0:03:50 EH: And so last year, we estimated that GiveDirectly, one our top charities, had about \$30 million worth of room for more funding, and the other three collectively had about \$17 million worth of room for more funding. And what that meant is that, we expected the donors that we were influencing to come pretty close to filling the gaps at the charities that were not GiveDirectly. And this year, we see significantly more room for more funding at these organizations. For the three that are not GiveDirectly, the Against Malaria Foundation which distributes bed nets, the Schistosomiasis Control Initiative, and the Deworm the World Initiative, which run deworming programs that treat children for parasitic infections in Africa. Collectively, those three organizations now have more than a \$100 million worth of room for more funding, and we think GiveDirectly also has about \$100 million worth of room for more funding. And the reason I'm separating them out into these two groups is that, what GiveDirectly does is direct cash transfers, which means they take money from the donor, and aim to deliver just about 90% of it directly into the hands

of very poor people in Africa. And so that is an intervention that scales relatively easily, and is able to take in a great deal of funding.

0:05:08 EH: So they've successfully expanded, and it's the other three that have had more issues with capacity over the last few years. So then the main drivers of this expansion on the GiveDirectly side, the cash transfer organization, they have successfully grown and demonstrated that they're able to scale over the last couple of years. When we first recommended them in 2012, their annual budget was around \$500,000 a year. Now, they're operating at a rate that's moving \$22 million a year into the hands of very poor people, in Kenya, in Uganda. And they think they could relatively easily scale up from there if necessary, or if they have the funding to do so. And based on their track recording, we're confident that they would be able to take in more money and use it well.

0:06:00 EH: With the other organizations, it's more of a situation where they've done more work and had good results, and those results give us more confidence that they could take in more money and use it effectively. So in the case of The Against Malaria Foundation, this is the group that funds distributions of insecticide-treated nets that prevent malaria. A few years ago, they scaled up from being a very small organization that did these 10 or 20,000 dollar distributions with very small charities in Africa to a group that was doing distributions that were about a million dollars a piece. And then due to GiveWell's recommendation, they then received enough funding to go to the \$10 million level. And when they reached that level they struggled to reach deals with African governments and ministries of health to fund bed nets; and we've written a lot about this on our website.

0:06:57 EH: And so two years ago, we actually took them down from our recommended list, because we were not confident that they'd be able to find distributions to fund. But over the past year they... So at the end of last year they successfully identified another distribution and they funded it and the main concern we had at the end of 2014 was the question of whether this second distribution would actually go well. They were working with a new partner in the Democratic Republic of the Congo, a challenging area to work in Africa, and so we wanted to see that they'd be able to run a high-quality distribution with this partner. And while we're still waiting for some information to come in, all signs point to the idea that they've been successful in this distribution in getting nets out the door and bringing them to people who need them and so at this point we would like to see The Against Malaria Foundation in a position where they are able to negotiate with as many governments come to them and are ready to... Who are looking for funding for bed nets, and that they have the funding to fill those gaps. And we think the total gap that they have before they would be bottlenecked by funding is around \$50 million.

0:08:10 EH: And then the other news is on the side of The Deworm the World Initiative. So this is one of the two groups that runs these deworming programs to treat parasitic infections in school children. At the end of last year we had seen limited monitoring information from them. They had only worked in Kenya and in a couple of Indian states and we were hesitant to recommend significant funding to them, not only because of their track record, but they weren't even asking for a lot of money. They were saying, "We don't think we have the capacity to scale up our programming." And over the past year, they've successfully expanded to three more Indian states. They've shared significantly more monitoring and evaluation information with us that meets the quality among the best monitoring and evaluation information we've seen from our top charities. They believe that they can use more than \$10 million to continue to scale up and reach more people with their programming, and they've shown some ability to do that.

0:09:09 EH: One of the exciting things that they were involved in last year is a national deworming day in India where they played some role in convincing the Indian government to have deworming days where all the children who were in [0:09:22] ____ school received the pills across five states in India that they weren't even working in. And so they've really had some success and this is an organization that we think now could take more funds and grow and we'd really like to see them have opportunity to do so.

0:09:38 EH: The other sort of major update and our recommendations this year is more internal to GiveWell, and so it's the way we're thinking about the criteria we use to assess and rank charities. So one of the updates this year is that we are putting more weight on cost-effectiveness analysis in our recommendations. In the past we've written a lot about cost-effectiveness analysis where we try to estimate the amount of life improvement or the number of lives saved per dollar donated. These estimates we think are necessary but also are subject to a number of highly uncertain assumptions.

0:10:19 EH: And so we've been hesitant to put too much weight on them, but as we've learned about these interventions over the years, and as we've been more intense about the analyses that we're doing, we've really come to the conclusion that there's really a case that some of these interventions are significantly more cost-effective than others. And so in this

case we think that it's meaningful to say that distributing bed nets or treating children for deworming is something like five to 10 times as cost-effective as cash transfers and that has major implications to how we rank charities. And so this year we are putting significantly more weight on that and really trying to see the charities running those priority programs, other things equal, get significant funding.

0:11:09 EH: The other big update internally is in how we're thinking about this concept of room for more funding. And this is really driven. The reason we had to think more about this this year was driven by two causes. One, I've already mentioned, the funding gaps themselves are much larger and so we don't think we're going to come close to closing them in the way that we were close to last year. And so that means that we had to do more thinking about how to prioritize between these gaps. Also, we're recommending that Good Ventures, which is a foundation that we work closely with and has supported our recommendation significantly in the past, we've recommended to them that they make significantly larger gifts to our top charities. Cari from Good Ventures is here, so, shout out, and so we because of that combined effect, the larger gaps and the larger expected funding from Good Ventures, and also from the individuals who use our research. We wanted to make sure that we were... We wanted to be more thoughtful about how we assessed room for more funding.

0:12:15 EH: In the past we had mostly just assigned a single figure to each charity. So, we would say something like, "GiveDirectly's room for more funding is \$20 million." But, we were not specific about whether that meant \$20 million that they would spend in the next year, or that they would need in the next year to plan for future years. We didn't talk about the likelihood that that would be sufficient funding or insufficient funding. And so, we've tried to make two changes to this analysis this year.

0:12:46 EH: The first change is that we think different types of funding can have different impacts on organizations. And so, we're thinking about three types of funding. One, is we just call incentives funding. We think it's valuable for the charities that make it to GiveWell's top charity list to be rewarded for doing so, because it creates incentives for groups to try and follow norms of transparency, to demonstrate that they have effectiveness. In many ways, this was a driving cause, or driving part of GiveWell's mission to try and create better incentives for charities to demonstrate effectiveness. And so we want to make sure that that incentive is there. And so, we're setting that incentive now at a million dollars. And so, for all the charities that make our top charity list, we recommend to Good Ventures to give at least a million dollars to them to make sure that that gap is filled.

0:13:40 EH: The second type of funding gap that we've been thinking about is what we call capacity building. And so the idea here is that sometimes an organization is operating at scale, and the impact of it receiving additional funds is that it's just going to move additional funds out the door to implement deworming programs or buy bed nets. But for some organizations, additional funds have a dual benefit of enabling them to determine that they... To learn how to operate at greater scale. So that could mean, in the case of AMF or Deworm the World, entering new areas that they haven't worked in before and figuring out the challenges inherent there, so that they then know that they can take in more room for more funding... They have more room for more funding in the future. In other cases, it could be providing them funding to build out a fundraising department, so they can raise funds from donors who are outside the GiveWell sphere of influence. And so it's a funding type that we see as highly valuable. And for our top charities, even more valuable than pure execution.

0:14:48 EH: And then finally, there's pure execution. And that's just where you imagine that the money's going in and going out to purchase more of the program that the charity is running. And there this year we've tried to basically model that as well as we can probabilistically. So, there's a room for more funding level, where if the charity does not receive that amount of funding, we think it is more likely than not that their progress will be slowed in the next 12 months, due to lack of funding. So, that would be a case, in the example of AMF, they have to negotiate with governments. There are times where they have, in the past 12 months, paused on negotiations, because they no longer felt confident that they would have the resources available to sign those deals were the other party to be willing to move forward. And so one level of room for more funding is just the straight assessment of how likely it is that they end up in that position.

0:15:39 EH: And then we've also tried to set, what we've called level two and level three levels of room for more funding, where, at some level, there's only a 20% chance that they would be bottlenecked by funding, and then at another level, there's only a 5% chance that they'd be bottlenecked by funding. And then we take all of these... We looked at each charity. The type of funding gap and the level of the execution gap, and are trying to rank them against each other in order to make our recommendations to donors. And so, when we end up publishing our recommendations within the next few days, what you'll see is this list of funding tranches for each charity, which we've prioritized based on program effectiveness, organizational effectiveness, and then the type of funding that we're talking about providing.

0:16:29 EH: So, I know it's a lot, but I want to pause there and just see if anyone has questions on this or really anything else related to GiveWell and our top charities.

0:16:38 HK: Let me just jump in with one thing about how the recommendations are different this year and this is kind of summarizing what Elie said, but last year, we said, "We want to get this much to this charity, that much to that charity, that much to that charity." And we had our recommended allocation, then we'd tell people to divide up their donations so we can try and hit those targets. This year it looks really different. It's, instead of list of four charities, we have a list of 10, or something, funding tranches. So it's first we want to get a million here. Then we want to get five million there. Then we want to get 10 million there, and they're in order of priority. Because sometimes we'll have a worse charity with a more pressing funding need, and then it will be followed by the better charity with the less pressing funding need, and then the worst charity again with the less pressing funding need.

0:17:21 HK: So we have this list of things, and then further complicating matters is the big... Grants we're going to recommend to Good Ventures, which are going to be very... A lot of money. And we're going to recommend that Good Ventures just goes right down the tranches and fills them. And so, first we have this list of 10 funding gaps. Then we're going to recommend that Good Ventures close a bunch of those gaps. And then we're going to take what's left and we're going to draw our recommendations to everyone else from that. And so how that's going to end up looking is, there's this one really big funding gap in that table that's kind of going to eat the end of what we're recommending to Good Ventures and have a lot left for individuals. And that's going to be an AMF funding gap.

0:18:00 HK: And so AMF is going to be our number one charity not because it's the best organization, but because after the grants are recommending the good ventures are done, that's where we're going to see the best combination of a pressing funding need in a strong organization. So it's under the hood more complicated, the bottomline is going to be a rank list of four charities and not a recommended allocation. Though we will put all those details in the blog post. Yeah.

[chuckle]

0:18:29 Speaker 3: It feels like in the past when I've heard you talked about room for more funding, it sounds more like a maximum than a minimum, like above this point we're not confident that they could spend as well. That sounds like most of what you were saying, you can sort of orient towards maximum but it was more of a minimum of, they need this amount to get to this point. Is that a change in approach, or is that still sort of the same thing in different words?

0:18:49 EH: Yeah, so I'm going to repeat the question just for this recording. So the question is, it seems like in the past we're often talking about room for more funding as a maximum that a charity could receive and now we're also talking about minimums. I mean, I think the answer is that in the past we were not thinking about this as thoughtfully as we are now. And so when we said, "This is the maximum a charity could receive," we were not saying, "This is the 50% likelihood maximum or the 80% likelihood maximum or the 95% likelihood maximum. And so I think that in that way it's just in a deeper way of thinking about the same gaps.

0:19:31 EH: We also didn't see the same types of minimum gaps in our top charities in recent years, because in recent years we've been more in a situation where charities were... We were more worried about the maximum side of the equation than the minimum side of the equation, because we had relatively few charities that were all relatively small and it seemed like it was more likely that we would overwhelm the maximum than not hit the minimum.

0:20:00 HK: We have done minimums in the past. I think last year we had a minimum and a maximum for at least some of the charities, where we said, "We'd really like to see them raise at least this much 'cause that..." and I think it was kind of an early version of the capacity building idea. And then we would say, "If they get to that much then we're going to think they're not even going to spend that money if it goes over that." So we've done stuff like that before but I think this year we just got much more precise about it and there's actually like now five different classifications of funding gaps and then they're all getting ranked against each other. The list, the ranking incorporates both how good the program is and how pressing the need is.

0:20:31 EH: Yeah, and I think one of the issues comes from the idea that if you talk about room for more funding as checking whether the charity is going to take in too much, that is sometimes how you might describe the concept in one sentence. But really, what the room for more funding analysis is about is trying to assess the impact of different levels of additional funding and that doesn't lend itself as well to just thinking about purely a maximum or a minimum.

0:21:02 Speaker 4: You mentioned [0:21:03] ____ a great effect is still on cost-effectiveness. Can you maybe talk about how some of the inputs into that work, especially when you're comparing sort of saving lives versus improving lives, and how you think about those two?

0:21:18 EH: Yeah, so how do we think about cost-effectiveness analysis especially given the fact that there's like comparing apples to oranges in many ways. So first off, we are totally onboard with the idea that deciding where to give to charity is not a math equation that can be solved perfectly. And so that the cost-effectiveness analysis that we do has to take into consideration people's subjective judgments about the extent to which they want to improve lives versus save lives. But our goal is to try and make those assumptions explicit and model them and then enable anyone who wants to to put in their own numbers and get the output that they find useful in making their decision about where to give.

0:22:07 EH: And so, to build the model we try to put the impacts that the organizations have into as similar terms as possible. And so with deworming programs, the headline benefit that we focus on is a study that found that 10 years after children were dewormed, they were earning 25% more than the controlled group in the original experiment. And so, deworming has this earnings benefit. And then similarly, a cash transfers program has an earnings benefit because it gives people money, they can spend it, they can also invest it. And so we try to model those things out so we can compare them against each other. And then comparing deworming programs to bed nets, this is a case where there's a sort of a standard way that the global health community tries to make these comparisons using a metric called the DALY, which is Disability Adjusted Life Year, where they try to compare... They put a weight of disability on a given condition and then say, "This is how it compares to being, basically a year of that disability versus the year of death."

0:23:12 EH: We don't just rely solely on this exact values that the global health community has because they're so judgment-laden and we don't find ourselves really... Well, I don't find myself at least, in finding those values very intuitive. So the value would be something like being blind is 60% as bad as being dead, and that... I mean, someone somewhere...

0:23:34 HK: That's an actual number.

0:23:36 EH: That's like the weight that they use. And so if you rely in the DALY framework explicitly you're just taking those numbers. Now, we are still more or less using that type of approach by trying to say, "In one case you have bed nets. The primary benefit of a bed net to prevent malaria is saving a child's life; the primary benefit of deworming program is trying to improve a child's life. And what you'll see in this cost-effectiveness model, which is on our website and you can play around with, is different staff members, their judgments about how they would make that trade-off." And so that's the process we go through, and so over the last couple of weeks, what's been happening here is everyone's been playing around with that model, asking questions, and trying to figure out what values they would put in. And I would say there's reasonable consistency across people in finding the bed nets and deworming significantly more cost-effective than the cash transfers.

0:24:30 HK: A couple things I want to add to that. One is it's, a lot of these things are hard to compare to each other, but we kind of do pairwise comparisons. So I think you can kind of compare cash and deworming because both of them, the main benefit, you're actually going to measure in dollars. And you have... The dollars are kind of happening at different times and in different ways, so it does get complicated, but you can kind of say, "Well, would you rather deworm this many children and have them get that much money later with such and such a probability and such and such an adjustment, or would you rather give some families cash and have them have this much money now with a bunch of other adjustments?" So then you're comparing dollars to dollars.

0:25:02 HK: Then, you can compare deworming to bed nets, where you can kind of say, "Would you rather have, save a life and get this many years of life saved, or would you rather get that many years of people earning more money because they've been dewormed?" You can compare those two. And so we kind of do it that way. I also think it's something that's worth just going over a little, like a mini mini history of just how we've treated these numbers, 'cause I think it's interesting and I think it also is important to what we're deciding this year.

0:25:29 HK: When we first started GiveWell, I think we were very excited about the idea of quantifying as much as we could and saying, "Let's figure out how much good each charity does per dollar." That's all we really need to know. We don't need to know, in theory, if that was what you knew, you could forget all the other information about how much they spend on overhead, how efficient they are, how good an organization, et cetera. And then we pretty quickly found that these numbers were not reliable. They weren't robust, there were major error bars, the estimates were very fuzzy,

we were making everything up. And so then I think we had this long stretch where our cost-effectiveness model got more and more complicated, harder and harder to understand, harder and harder really for anyone outside of GiveWell to make any sense of and really, most people inside of GiveWell.

0:26:15 HK: And so I think as of a year or two ago, we had this spreadsheet and it was just, "Well, what do you think about how you should count the benefits to dependents and what do you think about what you should use for the discount rate?" And there was just like a million inputs that we don't really have data about that you have to kind of guess at. And there were maybe very few people even on staff who really knew what was going on and understood this spreadsheet. And so then, in the end, we'd play with this spreadsheet and we would say, "Okay, well sometimes I get that deworming is like a 1,000 times as good as cash, and sometimes I get that it's like 1/1,000th as good, and my best guess is that it's like 3x as good." I'm not using this for much. And everything kind of feels very close together and it all feels kind of made up. And I think we're kind of crawling out of the hole now, in a sense, 'cause we've had this spreadsheet around for a while and we've been poking at and improving it for a while, and this year we got a lot more staff involved. And so we were really encouraging, and other people were jumping in. People on staff were jumping in to say, "I want to put my values in the spreadsheet, understand what's going on." And we had a lot of conversations about just what everything in the spreadsheet means and how you can fill it in correctly.

0:27:18 HK: And so, this year, I feel a lot less like it's this weird black box that three people understand, and a lot more that it's a complicated thing and not very robust, but a lot of people who looked at it; I've been thinking about it for years, personally. I've thought about it really hard. And the numbers move around a lot, but certain qualitative phenomena hold. So, deworming and bed nets being better than cash by a multiple has held for a while and holds across a lot of people and a lot of assumptions and so... I've been starting to believe these numbers again and they're starting to affect our recommendations more.

0:27:49 EH: Yeah, though, I think it's important to know, if you are... If you have the view that we had a few years ago, if you say, "I don't trust this model at all," then, you probably should give to GiveDirectly. They're the organization that we think is strongest, and they get downgraded by our assessment of their cost-effectiveness. If you say that, "I really value lives saved significantly more than lives improved," you should give to AMF because they're the organization that's saving lives. The other three are not saving lives at all. They're improving lives. And then, certainly our recommendations this year, the organization that we've had the most trouble communicating with is the Schistosomiasis Control initiative, or SCI. They're the group that we don't quite know everything that they're doing, we've struggled to even find out exactly how they're spending money, but they run this great program, and they've shared some monitoring information, so they make it on our list. But I think if you are skeptical that that matters. If you say, "This group is running a simple program. It's probably going fine." They're the group that's getting downgraded the most in our rankings because of GiveWell's struggles to get the answers it's looking for.

0:29:04 HK: I think it is a really interesting question just to think about SCI versus GiveDirectly, 'cause you've got one organization that we think is really outstanding, and then you have another one that we think, we've really had problems understanding what they're doing and we still don't really understand how they're spending their money. And we think that's a bad sign at this point 'cause of how much time we've put in. And we've had other issues too with their monitoring sort of being in our view, giving the wrong picture and not being kind of forthright.

0:29:27 HK: And so you compare these two organizations. One argument would be, "Look, deworming consistently looks 5X better than cash or more, and it's been that way for a while and it looks that way across a lot of staff members, and it looks real, and it looks consistent. So, even if GiveDirectly is twice as effective, as long as SCI is decent, it seems pretty unlikely that it's really washing out and it seems like you really want to give to the better program as long as the execution's decent. As long as it's like 50%; they're getting 50% coverage of their deworming than they would... That would kind of by the math, that would make them be doing more good than cash." And then the other perspective would be like, "Well, you're giving to a worse organization over a better one because of what a spreadsheet said. That's insane." I would have a lot of sympathy with both perspectives, and I found it really, really tough to think about that question. But I don't know. This year I feel like the numbers feel real enough that the first perspective looks stronger than it has in the past.

0:30:28 Speaker 5: In the past you've been, I think, somewhat modest about trying to self-promote GiveWell and drive more eyeballs to recommendations. And in the past the funding gaps have been smaller. So now that there's a lot more room for funding and you think that your typical traffic isn't going to close them, are you thinking about more self-promotion, more ways to convince more people to give according to [0:30:57] ____.

0:30:58 EH: Yes, so the question is, "Are we planning to do more on the outreach front relative to what we've done in the past." So we've made a couple of steps in that direction. We hired our first outreach associate, Catherine, who is now, she and another research analyst, Sean, are leading our outreach efforts. And so we are putting more dedicated time into trying to reach out to donors because we now have the capacity to do so. We're still definitely prioritizing research over outreach. That's for a couple of reasons. One is, we still see a lot of research to do. Primarily, the biggest bulk of that is on the open-end philanthropy side which Holden is going to talk about later. And so a lot of the growth in staff and the senior capacity is going to try and build that, essentially, that organization from scratch. And so that prevents us from being able to put the energy into outreach that we might if not for the existence of Open Phil. The second fact is that our growth has still remained strong. Our best guess now is that, excluding Good Ventures, we're going to move \$25 to \$28 million to top charities this year compared to something like \$12 million last year.

0:32:18 HK: Total for the calendar year, right?

0:32:19 EH: Total for the calendar year. It's still, even now we're doubling off a relatively substantial base and that is also giving us some confidence that we feel comfortable with not rushing to put more of a priority on outreach. And then finally, we've just always had this experience that the biggest owners find us and learn... They value the intensity of the research and the quality of the research, and so we want that to continue to exist. And I think that there are definitely both promising opportunities on the research front on the GiveWell side and also ways in which the product is not all it can be. This is now the fourth year in a row or something, where our recommendations are basically deworming, cash transfers and bed nets. And I think to some extent those are there because they're outstanding interventions, but there are surgical programs, immunization programs, water programs, education programs. There are things that we have definitely not exhausted in the way that I would like to, and I think not doing so becomes more of a problem as time goes on.

0:33:29 Speaker 6: My question has to do with that. You say that the number of the charities that you're going to be recommending year is four, same as last year. I'm not sure what they were before that. So my question has to do with your scope in terms of what you look at. You're saying that these are the four top ones, but what are you looking at? How broad is your approach in terms of different charities across the board? Are you focusing only on these four? There's gotta be a lot more that are out there. And how are you going to proceed?

0:34:06 EH: Yeah, so the question is, "There's only four top charities. What's the scope of GiveWell's research?" Conceptually, GiveWell is open to any charity that meets our criteria, but our criteria are ones that most... The vast majority of charitable organizations do not meet. So the criteria being that there's a strong evidence base that supports the fact that their programs are effective, that the programs are highly cost-effective and competitive with the charities we've already found, and then that they can utilize additional funding in an impactful way. When GiveWell started we looked at both charities working in the US and charities working abroad, and it took us about a year but looking at the outputs from those programs, it was clear that a dollar goes further overseas. You're just going to get so much more for the dollars you gave when you gave internationally.

0:35:00 EH: And then we spent years canvassing hundreds of charitable websites looking at their programs, looking for monitoring information. And in most cases... In all but one case the last time we did one of these big searches of websites, charities just don't publish... They don't have, and it's understandable because it's expensive, high-quality monitoring of their program, where they can convince you that their program is making the difference and causing the impact to happen in a way that it otherwise would not have. And so mostly, we end up relying on, and I think this is reasonable, on academic evidence that assesses the evidence for particular programs. And so with bed nets, we're not looking for the Against Malaria Foundation to prove that bed nets save lives. That's a fact that's been evaluated and demonstrated in independent trials of bed nets, and what AMF has to show us is that it gets bed nets to people and they actually use them.

0:35:56 EH: And so this then becomes a very high standard that many organizations can't meet, and that's okay. I think when GiveWell started, there was a possible path we would take, which would be lots of charities reviewed and recommended in different categories. And instead, we've taken a very different approach, which is to... Basically try to find the places that we... When we started, Holden and me, but now other staff, would actually give our own charity at the end of the year. And for that reason, we're very happy having a small group that we recommend because there's just not a lot of value for us in having the 15th ranked organization that is somewhat different.

0:36:38 HK: Yeah. In some ways it's kind of boring to have the same four charities every year, and certainly, we'd like more options and more fallbacks if something goes wrong with some of them. But in some ways it's like, this year it's

better news that we have the same four charities, they have massively more capacity. That's better news from our perspective, our mission, than if we had found 10 new top charities and the total capacity was less. 'Cause our mission was not to rate every charity in the world, our mission is to do the most with every dollar. And that is not unlikely to mean there's a very small number of places that are actually doing the best work, as far as what we're able to identify, as far as what we're able to know, as far as what can be verified and communicated to individual donors. And so, it is definitely a feature of GiveWell, and it's by design. We aren't aiming to have as many recommended charities as possible. We're aiming to do the most with the dollars, and have as much room for those dollars.

0:37:32 S7: Do you want to talk though, about the work that you are doing to expand, that you've done the last year to expand the list, and that we're doing to expand the list going forward?

0:37:44 EH: Sure. Is that the same as...

0:37:46 S6: That's what I was getting at. You got these four, but what kind of processes you have to look at some new up-and-comers who might be doing this sort of thing. Because you're encouraging that kind of approach, and so people are going to be responding... In general...

0:38:00 EH: Yep.

0:38:01 S6: People are responding to that approach across the board out there.

0:38:04 EH: Yeah. So, there's basically two things that are happening now. One is, we are actively looking... We're assessing the evidence for programs that we haven't assessed the evidence for before. Or it was a long time ago, and we're saying... So there's claims that charities can perform a cataract surgery, and prevent blindness for \$40, or immunizations, nutrition, education. So for a long time, we were severely capacity-constrained where we just didn't have the ability to do it. Now, I think, really over the last 12 months, we've grown staff to a point where I'm much more optimistic about our ability to look at those additional programs in the next year. And so I hope we make more progress there. The other thing that we're doing is actively saying to people, "If you want to start a charity that you think will meet our criteria, and you can convince us that you have a shot, we will seriously consider providing you funding." Or, "If you run a charity and you need help getting to the point where you would meet our criteria, we would seriously consider providing you funding."

0:39:07 EH: And so we've now done that in one case; there's a group that's implementing a conditional cash transfer program in Nigeria trying to... The condition being that HIV-positive pregnant mothers come into the clinic every three months to pick up the pills that prevent mother-to-child transmission of HIV on... So this is a startup group. They are definitely not one that is currently meeting GiveWell's criteria. But because of the program they run, because of their commitment to transparency and evaluation, we're optimistic that they could potentially be a top charity in the future. And then similarly, we've provided some funding, or Good Ventures are providing some funding but with this goal in mind of building the evidence base out for interventions.

0:39:50 EH: We rely very heavily on randomized control trials to assess whether programs are effective. And so we've provided funding to a trial that's trying to assess whether a small incentive to have families bring their child to be immunized increases uptake of immunization and increases those rates. And then finally, we're working with a third group called IDinsight, which conducts evaluations to see whether we can, essentially, partner with them on evaluating charities' programs, when the program seems effective but we have unanswered questions for which the charities don't have data. We want IDinsight to be able to go in and collect that data for us. And so we're trying to work on both of those fronts, follow the ones we have, look for more of the same types of programs that we've recommended in the past and then provide some funding to some startup organizations.

0:40:42 HK: And we do put a lot of work into this, and I am really excited about it, but I think it's worth noting just how slow the road to top charity can be. Especially, if you're... There's different stages Elie talked about. So there's an organization that are already doing a good program, the evidence is already there, we need to scale the organization. Then there's cases where it's like you need to build the evidence base; then you're going to need to find an organization; then you're going to need to scale it. And even once you have a top charity... AMF was running a million dollars a year when we first recommended them, and it's been kind of a long complicated road to get into a point that we're comfortable recommending tens of millions of dollars there. So, I think this is just... I think it's going to be the same four a while, probably. I think we're putting a lot of effort into finding more. We're looking at everyone else who might be ready now and we're doing, I think, a lot of what we can do to help others get ready in the future. But the path from

an idea to an organization that has an evidence base, and has a capacity, and is able to take in all this money is just a long path. It could be years.

0:41:43 EH: Well, so I think we should actually pause. And another... I want to talk to you about another topic. Holden is going to talk about this a little bit. So, I think we mentioned that, we've recommended that Good Ventures makes significantly larger gifts to our top charities this year. And so I think it's worth just talking through a little bit what we recommended and how we came to that recommendation.

0:42:04 HK: Right, and we'll take questions on anything after that, but this is something that we anticipate being something people wonder about, so I do want to address it. 'Cause we do work... We work very closely with Good Ventures, we share this office. Open Philanthropy Project is what I spend my time on, and Cari and I both work on it, that's where our time goes, and we're partners in that. And so, one of the questions that we have been wrestling with, and that's we, that's me, Elie, Cari and others here, is, "How much should Good Ventures give to Top Charities?" And I think it's a pretty interesting question because Good Ventures is going to give away a lot of money over the course of Cari and Dustin's lifetimes.

0:42:47 HK: It's going to be... The peak giving is probably going to be in the range of hundreds of millions of dollars a year, for quite a while, and that's something they could be doing today, if we had the places to put the money, and it's something that eventually we do foresee happening. And so the question becomes... We're in this interesting situation, because Open Philanthropy, which I'm going to talk about in the next session, is this younger project that is looking at what are kind of higher-risk, higher-reward, longer-term, more complex things you can do as a major philanthropist to make the most out of your money. And we're looking at things like funding work that aims to help improve policy, funding work that helps fund scientific research that can cause real breakthroughs, funding work on reducing global catastrophic risks.

0:43:35 HK: These things are areas that take a lot of understanding, and as you'll hear, it's a lot of work to try and understand that stuff and come up with the best ways we can find to spend money in those very complex... And areas where we need to build up a lot of expertise and staff, so where we sit right now is we're not there yet. Where we sit right now is, we don't know as much as we need to know, for example, about how to fund great science, and the day when we do could be years from now, it could be a decade from now, could involve having a really built-out team of scientific advisors and scientific staff, and it could involve having great things to do that fund breakthrough science, and that could add up to like \$500 million a year and we could at that point say, "This is way better than our top charities." So that's a possibility. Another possibility is that 10 years from now we'll say, "No, we checked out science, we don't think anything in there's as good as our top charities." And so knowing how uncertain you are about that future, what do you do today when you have these giving opportunities sitting in front of you that I think are great, but I don't know how they're going to compare to these future things that we find.

0:44:36 HK: So that's been a real challenge, and I think the basic rules of thumb we're playing with this year... And part of the reason this has come up is 'cause the room for funding is so big, and because Good Ventures has gained a lot of confidence in the top charities process, so we feel like there's kind of no conceptual limit. If we and Good Ventures decided this is the best use of money, there's no conceptual reason, except for some that I'll get to, that we wouldn't want to close that whole funding gap. And so the question is, "What do you do?" And so some of the ideas we're playing with, and I would say these are ideas that we've been talking about very vigorously for a week or two, but there's a lot of evolving to do. One is just to think about what percentage of the total capital do you want to give every year, and there's a kind of a private internal burn-down model that says, "If you want to give it all away in your lifetime," which Cari and Dustin do, "Then here's a way to do it," which kind of goes up, and then stays at a peak for a while, and then goes down, without making any too-sudden moves.

0:45:38 HK: And I would say we're at a very early stage in Open Philanthropy, and so I want to be on the low end right now, but not on the ridiculously low end necessarily. So I would say, if we saw a good giving opportunities totaling to 5% of the capital, 5% is kind of the standard amount that one thinks you can give away, while in the long run keeping the same amount of capital, 'cause it's kind of a decent approximation to an investment return, it's also the legal requirement for what private foundations have to give away every year in order to avoid excise taxes, though that doesn't effect us because we're thinking conceptually about carrying Dustin's net worth, most of it's not sitting in a foundation. So 5% is like the low end, and I would say because we haven't built out our staff, because we don't have the knowledge we're going to have, we want to be on the low end, we don't want to be kind of giving away all the capital, or even at a burn-down rate yet, 'cause we don't know what we're going to find. But I would say if we saw great opportunities adding up to 5%, that would mean we could take all those great opportunities and still the capital would

remain about the same, and so that seems like a... That's where we put that kind of target.

0:46:49 HK: And then the next question is, "What counts as a great opportunity?" How good does a giving opportunity have to be before you say, "This is better than having to spend the money if I save it and spend it in the future. This is the thing to do now." We don't have an answer to that either. An idea we're playing with is this kind of benchmark of cash transfers. The idea is that GiveDirectly is... They have a lot of capacity and they could probably grow a lot more if they needed to, and they could take a large amount of money and get that to some of the poorest people in the world, and we think that's a really good deal in a sense, like in some ways we think of it is a... Like a 100X return on investment, in that the people who are getting GiveDirectly funds have about 1/100 the income of the average American, and so if you kind of think that the value of money to your life is kind of a function of how much you already have, you could think of that as a really good return.

0:47:38 HK: And because we think this is... I think there's a few ways that thinking about this as attractive. But one way to put it is just, this is a really kind of intuitively good giving opportunity, getting money to the people who have the lowest income. It's an opportunity that we think will be around for a while, it's an opportunity we think we'll be able to absorb a large amount of money, and so if we see things that aren't as good as that, there's really no reason to give to them. But if we see things that are a lot better than that, then they become a little bit more compelling.

0:48:09 HK: And so kind of the idea we're playing with is when we see something that we think is solidly and significantly better than that concept of giving cash to the lowest income people you can, and when the total giving is under the 5% budget because it's not being pushed out by better stuff, then we're going to go ahead and make the grant. And then we get into the caveat. So caveat one is, our time is way scarcer than money at this point. And so this rule only applies to giving opportunities that we have finished investigating and we have put in all the time, we've understood them, and now we're trying to decide whether to cut the check or not, or how big to make the recommendation.

0:48:50 HK: So that's an important thing. I think it would be easy to interpret what I'm saying is, any time you hear about something that might be better than cash transfers you have to check it out. That's not true. I have very limited time to check things out, and we're prioritizing checking the things out that we think are the best use of our time. But once we've checked something out, then this question applies. Caveat number two, is implicit commitments. So if you fund someone, you fund an organization and then the next year you decide, you change your calculations, you change your other options, you decide you don't want to fund them anymore, there can be kind of a very... It's usually really bad for organizations to have a really big drop-off in revenue very suddenly, and so you have to think about the consequences and what you're getting yourself into. And so one way of thinking about this is to say, "When we recommend a grant we have to think about what we would do next year if we didn't think it was high impact anymore, to kind of be a good citizen and not get that organization into trouble, and we have to include that in the implicit commitment. We have to include that in the amount of money in that accounts against the cap."

0:49:52 HK: So that's another caveat that we think about. A third caveat... And there's going to be total of four. A third caveat is this idea that when we're early... When we're learning about a space, a lot of times recommending grants is an important way to build knowledge and to build capacity internally, to build our own staff. So a good example of this, Open Philanthropy Project we've written a lot about this but, a lot of times the only way to explore what it would be like to work in criminal justice reform is to say, "We want to make some grants in criminal justice reform. Who's got something?" And then once you say you want to do it, you probably should do it, once you see things that are kind of in line with what you're looking for. So that's a concept we've written about a lot but the point is, if we're early in the space, if we're getting to know it, we're going to have a lower bar. And especially if we're really early and we're just trying to build staff, if a grant is helping us build the knowledge and the capacity to understand a space, then that benchmark is not going to apply or it could apply much less, kind of depends.

0:50:51 HK: And the final one is the really hardest one and the one that has probably been the biggest thing we've all been discussing very heatedly. And I don't really expect to close the book on it in everyone's minds today. But it's what I call coordination issues. So the issue is let's say you see an organization and you want to fund them, and you think that a \$10 million would be a great amount of money to fund them at, but there's other funders who could fund them too. And some of those other funders are looking at them at the same time you are, and you're having some conversations, all of those other funders are just kind of... You don't even know who they are but you think they exist. You think this organization can pull in money from others. I'm going to say that I think it is a bad practice and it's going to get us into a bad situation and not be the best use of our money, if we get in to the habit and get the reputation of always closing the full funding gap, of always being good for all the money that we want an organization to raise. 'Cause what that's going

to do over the long run is that as soon as someone hears that Open Philanthropy or Good Ventures thinks a certain idea is cool, all the other funders are going to run away from it at top speed.

0:51:56 HK: That's not a good way to get the most of your money, and that's not a good way to help make sure the organizations you're supporting are strong, are diversified, aren't totally reliant on you. And so I think we need to avoid those kind of dynamics. The problem is that you don't want to go too far in the other direction, and so you also don't want a world in which every funder's habit is to wait until they're sure they know what every other funder is doing, and then just come in and fund the last remaining need. I think if everyone's trying to do that, you're going to get a lot of waiting, you're going to need a lot of not granting, you're going to get a lot of funders not being open with each other and not sharing information. So I think that's a bad thing too. And I think this stuff gets legitimately tough. Like you have different funders who have different goals and each of them wishes the other would fund as much as they could, and then they could come in and fill the rest. And so it's like negotiation, it's like game theory; it's complicated and there's never going to be a perfect solution. And it's especially hard if you want to create a short term effect of getting good things to happen, of getting the good organization funded, of not spending too much money.

0:53:02 HK: But you also want to create a long-term effect of people doing the right thing, putting in their share and being honest with each other. So a lot of times if we're looking to grant and someone else is looking to grant, a lot of times the best thing to do in the short term is to kind of be cagey, not tell 'em what we're thinking, see what they do, then put in the rest. But in the long run building relationships with other funders where we're being open with each other is important, and so a lot of times it's like, we'll tell another funder what we're thinking, we're hoping they tell us what they're thinking and then we all are kind of acting in good faith and paying what feels like our fair share. So this stuff gets very tricky. A time when it gets especially tricky is when we're looking at top charities.

0:53:37 HK: You're looking at top charities and you say, "Look, there's a huge funding gap here, it's in the 10's of millions, hundreds of millions of dollars and the other funder that we're worried about coordinating with, is not some person we can call up on the phone. It's an amorphous blob of lots of individual donors, and we can't sit down with them collectively around the table and hatch something out and come to what we think is a fair agreement. How do we coordinate with them?" I'll briefly say I think there's like out of all the discussions we've had over the last week or so I basically heard three schools of thought. There'd been a very large number of specific proposals, a lot of them are very complicated, but I would say they all fit into one of three broad families of approach to this.

0:54:19 HK: So family one is what I'd call the funging family. That's "funging" comes with the term "fungibility". The idea is that you're always trying to make sure the charities meet their funding targets, and what you're kind of doing is you're letting individuals put in as much as they can, and you're covering the rest. And the advantage of this is that you're always getting in the short... It's a short term approach.

0:54:41 HK: You always get the right amount of money to the right places, and you don't pay more than you have to, so you look at... You can say, "Well individuals are going to give 20 million, and the total gap is 100, so we're going to recommend 80," and then you didn't spend more than you had to and you closed the gap, and it looks great. The problem is in the long run, and the problem is in the incentives it creates. Because in the long run when you take that kind of approach, when you say as a big funder you're always going to try and close that gap, what you're doing is you're creating the incentive for other donors to not give. And I think whether they understand that is a different question, but I think we do not want to be in that kind of equilibrium, we don't want to be in that kind of world where what's really going on is that every time you give to a top charity, we just take it into account and give less. And so what really happened is you kind of gave money to a giant funder to do other things. That's a dynamic we don't want, and we believe that if that dynamic happened bad things would happen, individuals would give less to top charities, the total pie would get smaller, so that's the funging family. I said the strength, I said the weakness.

0:55:42 HK: Family two is the matching family. It's going in completely the opposite direction. It's saying, "Instead of worry about how much every charity gets we're going to obsess over creating the right incentives." And so you can say, "We're going to match every dollar given to individuals. For every dollar they give, we'll give a dollar." Well this is great on incentives, this has no problems. Now every individual's incentive is to give, and in fact they get a great deal because they're kind of getting this leverage effect. The problem is on the opposite side. You're getting very suboptimal allocations this way. You're letting the amount you give be determined by the amount individuals give. You're letting the allocation across charities be determined by the choices individuals made, and you're going to leave a lot of good giving opportunities on the table in some years, and in other years you're just going to massively overfund a bunch of stuff, and leave some good stuff on the table.

0:56:30 HK: So that's the matching family, and then the final one is what I call the split family which is, I would say by far the least elegant, and just has kind of like tone-down strengths and weaknesses on both sides, which is you basically try and come up with a fair share, and you just give that much and you're done. And I think that's what... That is I think a fairly common strategy of major funders who don't know who else is around, and who else... What those people are thinking is they'll say, "You know what? We can cover half of this. We think that's fair. We think that we're the biggest supporter here, we're the biggest fan, we can cover half, and we're going to expect the others are going for the rest, and we're just going to fund half the gap and that's it, we're not doing more, we're not doing less, we're not matching, we're not funging."

0:57:13 HK: I think the advantage of this there's no incentive problems. You just picked a fair share, you put in what you thought, you kind of pretended everyone is doing the same thing, and you would get a good result if they did. So I think you don't have the incentive problems, and I think you have a lot more control over your own giving, so you can say, "This is the right share and the more funding gap there is, the more money we're going to give, and our money is not tied to individuals, we're not chasing around individuals and letting them decide what we do." The disadvantage is I think this just doesn't feel good, when you pick what you think is a fair share, and then you fund it, and then the other side doesn't come in, they're not good for it, they don't hold up their end and then you have this gap that's just sitting on the table. So it's kind of unattractive, but the thing I would say on the kind of game-theoretic aspect of this is that, if you're never willing to do that, if you're never willing to leave some of that gap hanging there, then you're going to be stuck in a world of funging. Because if you're always determined to get that gap closed, then in the long run you're creating incentives for everyone else to get out of there and let you close it.

0:58:16 HK: So very tough issue, we are going for now with our recommendation, we're going with split and we're going with 50% largely because, I think, aiming 50% is kind of a logical number in that we don't know in the end how big the individual clump is, but we think it could easily be as big or bigger than Good Ventures, but Good Ventures is very big. And it's also I think a number that just very clearly signals that we're not doing the funging thing, we're not looking at what individuals did and trying to get the right result this year, we're doing something that makes sense in kind of a more timeless way and that's a number that wasn't chosen based on what people are doing, it's chosen based on what we think is right.

0:58:56 HK: And so what we've done is we basically took the total funding gap that we think would be of interest to both Good Ventures and individuals, and we said, "We're going to recommend that Good Ventures fill half of that total gap." And then that's it, that's all there is, and then the rest is whatever happens happens. Again, now that we've had these discussions we may consider like 20 other schemes for next year, but in terms of something that I think doesn't create problematic incentives and is consistent with giving a lot of money and giving in the most effective way possible is the best we're able to do at the moment. And then some of the specifics of how that gets done can get a little complicated, but basically we've done that for the total top charities pie. And then as far as how the money get allocated we're just recommending it get allocated as well as we possibly can, so we're going right down the tranche list, so we're saying, "First fill the best gap, next fill the second best gap, next fill the third best gap, until you run out of that budget." And then you fill as much as you can and then the rest is up to individuals, so that's how we arrived at that number, and that ends up in... Going to be a recommendation of something like \$40 million this year to Good Ventures, and that's going to fill a lot of those funding gaps and leave a lot on the table as well.

1:00:10 EH: Cool. So, I want to open it up for questions on that, or on anything else that we talked about, or anything else that relates to GiveWell's top charities.

1:00:20 HK: And we can go over a little bit and shorten Open Phil a little bit, 'cause I think there's a lot to chew on here and Open Phil is going to be a little bit less of, by those standards, a little bit less dense.

1:00:33 S8: You speak a lot about funding gaps in finite terms, and there's a finite number of actions that actors can make in the game theory model and it's a discrete one-year timeline. But isn't it really more of a continuous distribution that

ke any actions that you make, you can change down the road as soon as you see people behaving to different incentives? And can you continue to change your strategy? Are you basically locking in that what you guys are doing now is what you're going to do, and that that's what you think is the best move right now and you're going to re-evaluate it later? 'Cause it seems like it makes sense on a one-year timeline, but you mention that funding gaps are increasing each year, each time you evaluate them the funding gap changes. And really on a daily basis, the funding gap's changing. As soon as they get more money, at some point they're going to have the ability to expand more and grow more. So, can you talk a little bit about that and how it might change in the future?

1:01:26 HK: Sure. How much flexibility are we giving ourselves to the future, have we locked ourselves in forever? I would say one of the biggest reasons that the split approach is most appealing to me, is because I feel like it doesn't lock us in. I feel like we can declare this approach this year and it's simple. A lot of the other ideas that we had as we were having these discussions had to do with, well we're going to do this this year and this next year and it depends on what individuals do, and then we engineer the right incentives, but we're still also early in this stuff and we're still figuring it all out. We want to be able to change our minds next year, and so very purposefully we picked a strategy that doesn't look like it's about trying to fudge people and get the right result this year. Picked a strategy that looks like it's our best guess at what's always the right thing to do and then next year, we can have that best guess be more informed and better and do something totally different. And that is the plan. So the plan is not to stick to this, and because we're not sticking to it, I think it wouldn't be very credible if I said, "Well, we decided the right split was 81.3%, which happens to be what fills the funding gaps." 'Cause we are going to change it next year. I mean, we might not. 50% might be perfect, or the best we can do.

1:02:31 S3: But when we announce our top charities grants, we do commit to not giving more...

1:02:36 HK: That's right.

1:02:36 S3: In the short term because we want that split to be real.

1:02:41 HK: Yeah, and that is why the recommendations go to individuals and Good Ventures at the same time. So we're going to publish... Good Ventures gets a heads-up about what they're going to be in and so do... But we're going to publish, we're recommending this here and we're recommending that there, because we could do things totally differently. We could see what happens in December, which is when almost all the individuals give, and then just fill the rest of the gap in February. But we're not trying to signal or take that kind of approach.

1:03:15 S9: This is a question about, why are the charities the same this year as last? We've looked at the age distribution of charities generally? That is to see, what's the turn on new charities entering into the field, if you look at where are all the charities' [1:03:31] ____ distribution of ages, maybe weighted by funds. What would be the base rate just from new charities coming in, that we might expect we can turnover, if you're already had [1:03:44] ____ effect of a perfect ranking of every charity?

1:03:49 EH: So, the question is whether we've looked at how frequently new charities come into existence and how old charities are generally. I think it's like... One way you could think about this is we've been sort of looking for the same type of charity for eight years, and so therefore you shouldn't expect existing charities to be a very good fit for our criteria, and so therefore the only charities that we could possibly evaluate would be new ones that came into existence. And it is true that all four of our top charities were created since 2003, and two of them were created in the last five years. And so, I think there may be some truth to this argument, though, another major change at GiveWell relative to where we were when we started, is that now that we are moving significantly more funding, we're able to get access to many more organizations than we were in the past. So, the first time that Holden and I tried to call a charity, they basically just laughed at us because we were asking way too many questions given the possible reward they could get for talking to us. And now we're in a position where, due to the amounts of funding that they're receiving, some of the charities that I look to next year as possible recommendations are ones that have existed for a very long time. And it's GiveWell that's changing in a way that makes them potentially promising as a top charity.

1:05:13 S10: Do you expect five years from now, that it'll be cheaper or more expensive to do the same amount of good? And maybe it's worth dividing that question into the direct good that GiveWell's top charities do, or the more quantifiable good that GiveWell's top charities do, and potential Open Philanthropy project candidates would do?

1:05:37 EH: Got it. Yeah. So, the question is, "How do we think the cost-effectiveness of charitable opportunities is changing over time?" Just on the GiveWell side, and you can answer for Open Philanthropy if you want, but I think that, my guess is that eventually, the opportunities will start to get a lot worse. Bed nets that prevent malaria, it's just a unbelievably good giving opportunity that exists right now. It costs... Now, nets have actually gotten cheaper recently. But while there's a big enough net gap worldwide, that we don't have enough funds to fill the full gap. I think the net gap right now is in the low hundreds of millions for the next couple of years. That is a gap that I would guess is going away because other programs in the past that have had similar levels of cost-effectiveness are ones that have been filled relative to where bed net sits today.

1:06:37 EH: The same is true of deworming. It's just unbelievably cheap to deworm children. There's a finite number of children who will need to be dewormed, general development makes it more likely that deworming programs won't be necessary. And so I think eventually we will get to the point where the opportunities don't look nearly as good as they do today. At least on the GiveWell side of things, and I think that's part of the argument for giving more now.

1:07:00 HK: Yeah, I think in the Open Phil side, I think the actual opportunities that exist in the world are getting worse, but the opportunities that we're able to see and identify are getting better, because I would say our knowledge of the world is improving, our staff is growing, our capacity is growing. And so I expect what am I able to be confident in today, what am I able to point to today, versus what can I point to in 10 years, I'm going to be way, way, way better in 10 years. But not because the actual opportunities will be better. I think the world will be wealthier and there will be more philanthropy, and so therefore the best opportunities will be worse, but the best opportunities we're able to find will be better. And I think those two things have tended to push in opposite directions. I think with GiveWell too... I think the opportunities this year might be the best they've ever been, because we've been recommending these charities for a while, and they used to be much smaller and have a lot of sputtering as they kind of grew, and now there's this big gap with these organization that have kind of learned how to do it and built up a track record.

1:07:58 HK: And so for a while it was like the capacity building on the charity side was going in the opposite direction. But now I think these opportunities are really good, I don't really see how they're going to get a lot better, and I definitely see how they're going to get worse, which is that hopefully, everyone in the world will be getting dewormed 10 years from now. It wouldn't take that much money for that to happen. And I really hope that happens, and that would make our options worse.

1:08:20 EH: At least everyone who needs it will get dewormed.

1:08:22 HK: Yeah, true! Fair! [laughter] I've heard that pills can cause nausea if you don't eat first.

1:08:30 S11: To what degree if at all are you concerned about the replicability of randomized controlled trials and other studies that you rely on? I think about this in the context of early childhood education, where you just see numbers all over the charts. Maybe your answer is, "Our interventions are so simple that there really is no replicability and scalability issue," but I could see that coming up.

1:08:49 EH: Yeah, so the question is, "How do we think about the replicability or the external validity of our randomized controlled trial to the program a charity's running." So in the case... You can imagine a charity runs one really great randomized controlled trial at one time in one location, they had some effect and now they scale up the program globally. Do you really expect those results to hold in other locations at other times? The short answer is we're extremely worried about it and it's a major focus of what we do. In the case of bed nets, it's not a single trial, there've been more than 20 trials of bed nets all over the world. And collectively they show that bed nets prevent cases of malaria, there've been five trials that explicitly measure child mortality, they show these significant effects on mortality. And so the evidence that bed nets prevent malaria and save lives is extremely strong. There's definitely still other questions that need to be asked about malaria transmission dynamics, about whether people really use the nets that they receive, and most importantly now I think the biggest worry relates to insecticide resistance, whether mosquitoes have built up resistance to the insecticide in the nets, such as the nets are significantly less effective.

1:10:01 EH: And this is something that we are worried about, we've written a lot about, it adjusts our cost-effectiveness model down slightly, but everything still points to nets being extremely effective. This is actually a much more salient issue with respect to deworming. It's a program where altogether, the evidence is much, much weaker that deworming actually has the types of impacts we're most excited about, but the program is so, so cheap, and the results that have been found, not just in one trial but in one trial and other studies that all of which have some serious methodological limitations, lend credence to the case that they could have these exceptionally strong benefits, and so from a cost-benefit perspective, deworming ends up seeming pretty good. And so it ends up being the case that this is a major part of our research.

1:10:54 EH: The incentives for immunization studies that I mentioned earlier, the funds that we provided, those are actually to two replications of an original study, and as part of our efforts to support the development of future top charities, we're funding additional data collection around a randomized controlled trial that was performed to encourage farmers to migrate during seasonal famine in Bangladesh with the understanding that the results that were measured in 2008 certainly might not replicate today, and so we're trying to get that additional data collected so that we can at least have a better view on how likely it is.

1:11:29 EH: Interestingly, just as part of that program, the charity went to two African countries. I think they went to Zambia and Malawi to see whether a seasonal migration program could work there. And after learning more about the economy they came away with the conclusion that it wouldn't in Zambia specifically. I think they just didn't have the same type of job opportunities existing in one part of the country at the same time that famine was existing in another part of the country. And so they were able to see ahead of time that the results weren't likely to replicate. That's a study that Evidence Action, the group that we're working with, they actually published on their website. So it's an interesting read for those of you for whom that sounds like a good evening.

1:12:11 HK: I'll add to that, this cost-effectiveness analysis we do with the spreadsheets, a lot of that turns into conversations about how good is this study really and what are odds it would replicate. So one of the things I like about the cost-effectiveness analysis is it really makes us think through all the things we've assumed without realizing it and debate each of those things and certainly, it's worth noting that with deworming, most people's calculation if you look at it... If you just imagined that the original study that was done was totally correct and totally representative and represented the universal truth of how deworming goes, then deworming would be 50X as good, as how people are modeling it. And so, instead of being 10X is good as cash transfers, it'd be 500X is good. So the implicit thing people are doing is taking that study at only 2% of face value, and I think there's a lot of reasons to only take it at 2% of face value. I think it's a really problematic study, and I think we've been just all over that study by now, and had just thought through every way in which it's so limited and we wouldn't expect it to generalize. So, I would say that we really do obsess over this stuff and the cost-effectiveness discussions are one of the times we obsess about it, although, not the only time.

1:13:18 HK: And I would say overall, deworming... Because it's a health intervention and because there are multiple data points pointing in similar directions, all of them have their own issues. I think deworming is like one of the very few most proven interventions that exist, but I would still say it's really far from proven, it needs like a massive discount.

1:13:35 S11: That's a pretty big gap. Doesn't that argue for replicating the study properly being a... How much are the problems of the study inherent versus just it wasn't done well? Maybe some fraction of the money should be going into investigating the study.

1:13:50 HK: Sure.

1:13:52 EH: The biggest issue with the study is that it happened at a time when worm rates were extremely high and so it's hard to know like how... So we know that the worm rates in that location at that time, because it was an El Nino year and there was flooding around Lake Victoria, children had way higher worm rates than they have in the programs that are now common around Africa and South East Asia, and so it's a major question. That right there, ends up being something like a discount to 10% if you're taking the... If you're just thinking about the study at face value. We have thought about this question of, "Should you go and try to replicate it?" and the... I think it's definitely... That question is definitely a live one.

1:14:34 EH: The biggest problem with replicating that study is that we wouldn't expect to see the results that we're most excited about for 10 years after the study completes. Because the result we're most interested in is the earnings benefit, you need more middle school children and so all of a sudden, you have to wait basically, 15 years from the time that you decide you want to fund the study until you get the results. And I think there's a decent chance that 15 years from now, everyone's been dewormed anyhow because it's so cheap. We're talking probably about... Well, you shouldn't quote on this number but I think the ballpark is \$100 million gets you very close to closing the global worm gap. And so it's just not clear that the investment in the trial is really worth it.

1:15:13 HK: Another issue that came up is because the original trial had such elevated worm rates, it was easy to detect and affect, and so if you believe the effect is only a 10th as big, you need way, way more people in the study to pick it up and then the study's just going to get to be a monster and it's going to be maybe prohibitively expensive or certainly, logistically, a huge challenge.

1:15:31 S12: In particular, if you want to detect an effect size that's 10% as big, you need a 100 times more people.

1:15:37 HK: And it was already a \$10 million study.

1:15:39 S12: So. Yeah.

1:15:40 HK: Yeah.

1:15:40 S12: We are thinking about how to do follow-ups to previous studies that have been done. It might be time for another follow-up to the same key study to look at five years later, what happens. I mean, another five years. So this is something we're very actively thinking about.

1:15:55 EH: Great. So let's use this pause to take a quick break, get more food, use the restroom and the circle back to talk to Holden more about Open Philanthropy.

1:16:05 HK: Cool.