

**GiveWell San Francisco Research Event, December 11, 2013**

*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.*

**0:00:02 Elie Hassenfeld:** So thanks everyone for coming. We really appreciate it. A big part of what GiveWell does is trying to make as much of our research public and engage in open discussion. And so we really appreciate you all coming to engage in this discussion with us.

**0:00:15 EH:** One thing you'll notice is we're recording this consistent with this idea of making as much public and we're going to post a recording and a transcript on our website after this. If there's anything that you say that you don't want to be in the recording or the transcript, feel free to say so while you're talking or just let us know after the fact. We're happy to remove it.

**0:00:32 EH:** The basic way that tonight's going to go is we're going to have this session for about an hour. It's going to be casual. If you need to leave, feel free to head out. I'm going to give... Talk a little bit about some of our traditional work on charities working in developing world and we'll take some questions. Holden will talk a little bit about some of the work we're doing on newer areas, GiveWell Labs, and we'll take some questions. And then we're happy to stick around afterwards to talk one on one with people as long as they want.

**0:01:00 EH:** Before we get started, I just wanted to quickly introduce all the GiveWell staff just so you know who they are. So if they could just quickly stand up to help me remember all of them. [laughter] We've grown a lot in the last year. So that's Tim, Sean, Alexander, Howie. I guess you can read, you might know this information too. This is Ben, Josh, Eliza, Jake, and there's Cari who is at Good Ventures, a foundation that we work really closely with and she shares our office here. She's also on our board. I was going to get there Alexander, always impatient.

**0:01:34 EH:** So what I want to talk about is some of the updates from our top charities. Our international aid groups that we've recommended in the past. And there are basically a few big updates on past top charities and some new top charity. So the first update, and this is something we've written about, you may have read it on our website, is about the Against Malaria Foundation. They distribute malaria nets to fight malaria in Africa. And this group has been our top-rated organization for the last two years. Donors have given about 10 million dollars to them as a result of our recommendation. And this year we decided to not recommend them again because they hadn't committed the bulk of the money that they had received to a distribution. The way that they work is they give money to other partners in the field who distribute nets. And they weren't able to find a way. They weren't able to use that money to make, to form a partnership. And so basically over the two years that we'd recommended them, they hadn't spent the bulk of their funds.

**0:02:34 EH:** So I want to go into a little bit what we think happened there. So firstly, the reason that we gave them our top ranking in the past is, malaria nets are a program that has a very high degree of evidence supporting it. There's a number of very rigorous studies that show that if you give out nets it prevents malaria and saves lives. It does so cost effectively. And AMF was particularly transparent in their approach. They would go and collect data confirming that the people who were supposed to receive nets actually did so. That the people... They surveyed people

about how often they used the nets. They tried to collect data showing that... Assessing the degree to which malaria fell in areas to which they distributed nets. So we liked all this about AMF. They stood out from all the other net organizations that we had been in contact with over the years for sharing this type of information with us.

**0:03:27 EH:** We first saw this issue of them not distributing their funds a year ago. And we talked to them and it seemed like we understood the issue, that they had a relatively short list of countries that they might work with. The countries fell through and they didn't finalize the distribution. And they seemed to address it earlier in 2013 when they had a larger bench, so to speak, of possible countries to work with. What we ended up seeing at the end of the year. And I should say that our change in our recommendation of AMF is really a function of a belief that, given the amount of money they have now, it makes less sense for donors to give more there rather than other places as opposed to a change in our assessment of the quality of the organization. Nonetheless, there are things that we think they could have done better. I think some of these things they disagree with us about. And you could read their post on the same topic on their website.

**0:04:20 EH:** But basically there are three main factors that I think drove the inability to finalize the nets distribution. Big picture, AMF was very successful at a small scale. It was distributing on the order of a million or two million dollars a year working with non-profits directly, and it was able to do that reasonably well. And through the funds that it received from GiveWell, it scaled up to be significantly larger. It, I think, currently has roughly 12 million dollars in the bank, and that change in scale changed the way that they worked. They went from working directly with smaller non-profits to negotiating with the governments, the ministries of health, and the big players in malaria control; the US Government, WHO, UNICEF, and the Global Fund.

**0:05:06 EH:** And this change in scale led to issues in their trying to finalize these distributions. So, one of the main issues that they faced is they have relatively stringent reporting requirements for the distributions that they support. Now this is something we support, better monitoring and evaluation data is something that I think is needed in malaria control. At the same time they were, they are, a relatively small player, relative to these bigger organizations, and they weren't in a position... Yeah, I think the other organizations said, "You're not, essentially, contributing enough to encourage us to agree to these more stringent reporting requirements." Another issue they faced was one of communication style that I think they may have been more... When they worked with smaller non-profits, they were able to say, "If you want the money, you have to follow these rules, you have to come through on our requirements."

**0:06:00 EH:** And as they started playing with bigger players, that same communication style didn't work as well. And then finally, they have a position that they want the funds they receive from the public only to go to specifically buy the nets themselves, that would go to countries as opposed to other costs like shipping and distribution costs of nets and monitoring evaluation. This is not the case in every single donation they receive, but unless specifically earmarked for it, they largely want to spend the funds the public donates just on that. And so, the bulk of the funds they received had come through in that way and this led them to, in our view, negotiate relatively hard with countries to not pay for these other costs. Now, they, AMF would have a different perspective on all of this. They would say, "We can still use money." We think that the specifics of each country situation, just to summarize, would have led to the problems finalizing distributions.

**0:07:00 EH:** Regardless, we think the best view to take with AMF is essentially to watch and wait.

If they finalize the distribution, we would want to reconsider them, but as long as they have the amount of money they have in the bank, we'd rather wait and see and find out what they are able to do. We do think that they are likely to address a lot of the concerns that are raised. They seem to have taken steps that would do so. They've involved their Malaria Advisory Board, which is a group of well-respected malaria scholars to help them identify distributions and improve on the types of challenges that I mentioned earlier. So, we're generally optimistic about the likelihood, that the situation with AMF improves in the future.

**0:07:42 EH:** Another update... So, that one was longer. The next three will be a little bit shorter. The next update is on the Schistosomiasis Control Initiative or SCI. They're another one of our top charities from last year, they remain recommended this year. They run a program that treats people for parasitic infections that are very common in the developing world. The program is known as de-worming and this is an organization where there's both some good and bad news. The good news is that we had always struggled to communicate with SCI relative to how well we communicated with our other top charities about what they were doing and I actually wrote a blogpost earlier this year talking about why I hadn't given to SCI last year even though they are one of our top charities, and it was because I didn't feel like we had enough visibility into their operations for me to feel confident giving my donations there. This year that's really improved, we have significantly better understanding of what the organization is doing relative to what we had before.

**0:08:38 EH:** The negative update on SCI is that a large part of the case for their programs were a series of monitoring papers that we thought showed that when SCI ran its program, kids started out having high levels of infections and then afterwards, these kids were measured again and the infections fell. And so, that was some demonstration that even though we didn't have a level of visibility that we wanted into the organization, here were some relatively strong documentation that they had run their program and been successful at achieving the goals that they set out to. We learned earlier this year that these studies weren't showing what we thought they were showing. We thought they were basically studies of kids who were treated through SCI's program and instead, it seems that these kids were initially tested as part of the program, but then if they were found to have the worms, immediately treated by the tester. And so, this clearly was showing something very different than SCI's general program which is just going out and treating everyone in a given area at mass. This was instead a paper trying to study the effects you could see from a program that is somewhat similar, giving a pill to a kid and following up with him a year later.

**0:09:55 EH:** It's not entirely clear to us why this study was necessary, given that there's a lot of studies that have shown the effects of treating a child with the pills that cure the worms, but nonetheless, the fact that this study doesn't show what we thought it did and we don't think or we haven't now seen evidence that SCI's programs are themselves effective has led us to, has taken away a big part of support for their program. As we were doing all this research and going back and forth with them, they sent us a few additional studies that may be the type that we were looking for before, we haven't had a chance to look at those yet. So, we're going to look at those and hopefully, we'll know more in the somewhat near future.

**0:10:39 EH:** Then, there's two other charities. One is GiveDirectly, they're another one of our top charities. They do something very simple. They just give cash to people in the developing world. Basically, for every \$1,000 they get, \$900 goes to a very poor family. The big update there, we've followed up with them over the year, everything still looks good. They're running their program, they're monitoring, the results are pretty good. But the big update is that a randomized control trial

of their program came out recently and the results were very strong. And this is big news for two reasons. First, it's the only charity that we know of, that's on our list, that has a randomized control trial of its program itself. And so, in some ways, GiveDirectly is arguably the most proven, most monitored charity of the ones that we recommend.

**0:11:27 EH:** Also, a big question that people had about GiveDirectly before is the case that cash transfer programs in general have been studied extensively by development economists, but GiveDirectly's program differed from the standard cash transfer program in a couple of ways. First, the transfers were much larger but also the program was unconditional, meaning they just gave cash to people without requiring the people to do anything to receive the cash. And most of the studies, not all but most, are focused on conditional cash programs. And so, this study was, to some extent, not only an answer of how GiveDirectly's programs worked specifically, but also a data point on the type of program that they were running for which there wasn't evidence that existed before.

**0:12:18 EH:** The final organization I want to talk about is Deworm the World, which is part of a larger organization called Evidence Action. This is also a deworming charity like SCI and they're our new top charity this year. They differ from SCI in just a couple of ways. Where SCI is primarily focused on using money to directly deworm children, Deworm the World is instead more of an advocacy and technical assistance organization. Which means that they are largely trying to encourage governments to run this program when they otherwise might not have and then assisting them in running the program appropriately. And so that raises one big challenge for us because the assessment of impact in that case and attributing impact to Deworm the World as opposed to other factors is very difficult. So that, for us, was a new type of endeavor this year, but we largely felt that they were an organization worth supporting.

**0:13:12 EH:** The other big difference between Deworm the World and SCI is that SCI works in Africa. Deworm the World works in India, and generally speaking, the degree of these infections in Africa is significantly worse than India. It's already somewhat unclear what life impact treating these worms has on people, but with the significantly reduced burden that Deworm the World is addressing in India, I think there is some question about the extent to which the deworming programs are having an impact on people's lives and that's something we're not entirely sure about. So, that's the quick update on our GiveWell top charities and so now I just wanted to open the floor for any questions about what I said or anything else related to our work on international aid stuff.

**0:13:58 Holden Karnofsky:** Just a real quick comment on questions. We know there's a ton of information on our website, we write like a million blogposts, they're really long. The purpose of these meetings is for people to not have to engage with all of that if they don't want to. So that's why we give this kind of summary, the highlights of what had happened this year in GiveWell top charities. And so, feel free to ask whatever questions you want and don't feel like you need to ask something that's as if you already read everything we put out because we don't really believe there are many people who've done that.

[laughter]

**0:14:25 EH:** Yeah? Right here.

**0:14:26 Speaker 3:** So this comes to the point of, will scaling be an issue going forward always? Like, all charities start out being limited in scope, doing something which is small, but then grow

and maybe addressing a small piece of a large puzzle. But then once you always scale, you need to add other costs and do other things, and maybe governments aren't designed for that kind of support from two other charities or whatever, whether it's tax or whether it is logistical support or whatever. So how do you think scaling will feature going into the future?

**0:14:57 EH:** Yeah. Good question. We're going to repeat the questions just for the sake of the recording 'cause they don't pick up audience questions so well. So the question is, how do we think about this issue of scale for charities? This seems like a common problem that many charities will face. The simple answer is, yes. I think it's like a giant problem that we need to try to address, and a lot of the work that we're trying to do with GiveWell in giving money and then following up with charities about how well it's spent is an attempt to try and learn about the challenges of scaling. I think we learned something with AMF in the challenges that they had scaling. We learned things previously with recommending an organization called VillageReach, where they also received a significant amount of funding and we said they similarly had enough, and we're trying to use what we learn to inform our decisions going forward with an organization like GiveDirectly.

**0:15:46 HK:** I would add that there's pros and cons to taking a charity to the next level of scale because scaling is such a challenge, if you could take a smaller organization that's doing something really good and it's doing it well, and you can help them get enough money to scale up, that's got major implications for the future. On the other hand, it's a lot more risky than just pouring in your money into a large organization that's already working everywhere and so I think there's pros and cons to that.

**0:16:13 Speaker 4:** So, it sounds like the randomized trial of GiveDirectly addressed a lot of concerns. What are some of the remaining concerns or possible weaknesses with the clinical trial, with the randomized trial?

**0:16:24 EH:** So, the question is, what are the remaining concerns with GiveDirectly after the clinical trials. So, I...

**0:16:29 S4:** Or about the trial itself.

**0:16:31 EH:** Yes. I'll give you a quick answer on what I think are some of the questions I have about GiveDirectly, and I think people I have not been... I think the people who did the work on looking at the trial should give us some of the major unanswered questions, and I think the biggest unanswered question about GiveDirectly is the question of scale. They are small and they are trying to grow, and it's not entirely clear what challenges they will run into as they try to do that and at what point they hit this inflection point where they can't just linearly scale up the thing that they've been doing, but start to do something that is different in kind, which is the problem that AMF ran into.

**0:17:04 EH:** One quick note on the trial. I mean it's still only one trial in one location. It also was a variation of GiveDirectly's program, to some extent, many of the grants were smaller, somewhat like half of the grants were, I think, smaller amounts than what GiveDirectly does right now. But, do you guys, Alexander and Howie, is there anything you guys would add?

**0:17:22 Speaker 5:** I mean I would briefly just add two smaller things. So, I think Elie's concern about scaling is right, it makes sense. But I'd say like the two questions that we don't... That are still

an open and important is like... Roughly, like, what are the investment returns. This is a really big questionable effect how you think about the benefits of cash. People who spend a lot of money like investing in their business or other things like that, and how good do you think those investments are will play a role in how good you think cash transfers are. We still don't have a great long-term answers to that question for GiveDirectly because the span of the RC team was short. And then in the second question. I think is what's the optimal size for cash transfers? I think you got a ways to go before that question is answered vigorously. The sample's size of the RC team was pretty small to try and address it.

**0:18:05 HK:** I'll add a couple of more things about the trial. One is, I think we would like to know more about what's not only the optimal size, the optimal framing of cash transfers. So under what period are they given out? Should there be conditions attached? Should there be kind of suggestions attached? And then, I know some people are concerned about people spending the money on alcohol and tobacco. So there've been a lot of studies on cash transfers, and overall, they have just not found evidence to support this concern, but that evidence does... I mean, it's self-reported. It's asking people what they spent on, and so I think there is value to it, but nobody's done the study yet that's kind of used some of the methods that are out there for getting around self-reporting bias on to address that question. People should know that that question has not been, conclusively addressed, although it has been addressed to some degree and I don't see it personally as the biggest concern.

**0:18:51 Speaker 6:** So with regard to scaling and GiveDirectly, they give to people based on proxies for poverty, right? So if people know like not having a metal roof or something, so doesn't... If people know about the existence of the program, I guess depending on how big it is relative to the potential target population, wouldn't there potentially be moral hazard problems or incentive effects, like, "I'm going to not have the metal roof 'cause then they'll think that I'm poor?" I guess what I'm saying is maybe scaling as a... You might identify, like, where the... A predictable effect of scaling, if they're well-known, right?

**0:19:29 EH:** Yeah. So the question is as GiveDirectly grows, and they become more well-known. Aren't there ways in which you can even predict the problems that they would face? So on this question specifically about the metal roofs and the thatched roofs, this isn't an area we've been experimenting with so there's cases where they're just giving out money to everyone. I mean, you visit these places, I did that a couple of years ago, you can also look at the data. I mean, everyone is incredibly poor. I don't think the difference is between like thatched and metal is all that big a deal in sort of like the big picture.

**0:19:58 EH:** I mean, I also think you can identify some specific issues with scale, whether it's... Like, when you go from giving a million dollars out in a country to, I think they hold far more, 50 million dollars, you all of a sudden expect turn into problems with like... Or shouldn't say problems, but you at least need to address local government, potential for crime, political issues where if you're a representative, you probably prefer that GiveDirectly serves your area over another one, like your constituents will be happier. I think there's all sorts of issues like that and these are the things I think we're sort of just trying to think through.

**0:20:34 S?:** The problem you mentioned with SCI's study seems like a pretty big one. Like what does it mean then that it's still a top charity? Like how confident are you that they're good?

**0:20:47 EH:** Yeah. So the question is this SCI problem seems like a really big one, how confident

are we? So I mean I think this is a tough question. The biggest, one of the biggest things we've struggled with over the years at GiveWell is, what level of documentation of the impact is reasonable to expect from a charity in order to be confident that they're having impact. So on one hand, you have these studies and you say, "We used to think they showed impact, now we don't think they do. Does that mean that should reduce your confidence in the organization?" and it does. On the other hand, Deworming is this incredibly simple program. SCI as an organization is, I think, reasonably reputable, they're run by a guy who has been at the forefront of deworming programs over the last 40 years. And so, what would you really bet is happening with SCI, like are they really failing to deworm kids or not?

**0:21:36 EH:** And I think it's still reasonable to make the argument that if you say just because they don't have this type of evidence, they're not effective; that seems a little bit too skeptical of what is likely to be occurring. It's worth remembering that in the case of... In the case of GiveDirectly, I think they have pretty good evidence that people are receiving the money and spending it. That would be pretty shocking if that weren't happening based on the data we've seen. Neither Deworm the World or AMF, the other two charities, has evidence that is, would be comparable to the evidence that SCI... We thought SCI had that they no longer have. This is not just something that I think that you can generally get from a charity and so a lot of the push for keeping SCI is the... The bet that this is a reasonably, reputable organization running a pretty simple program that is incredibly cost-effective, but we'll see, I mean I think if we looked at these studies they sent and they've had big problems too, like that would also lead us to question it.

**0:22:34 EH:** The biggest update I think, the biggest substantial update that these studies give us, more than the lack of impact, is a question about how well we're really able to communicate with this organization when, for years, we had thought studies meant X. They had presented them to us as implying X and they didn't demonstrate X at all. Yeah.

**0:22:55 S?:** I have a couple of really basic questions about GiveDirectly. And these are probably in the category of [0:22:59]        about that. By what metric do you evaluate the benefits of direct cash transfers as being of the same scale or of the same level as, like, curing worming or preventing malaria? And what or how does GiveDirectly's selection of which people to give money to factor into that benefit?

**0:23:25 EH:** Yeah.

**0:23:26 HK:** So the cost-effectiveness analysis we did... We try to compare... When we do our top charities analysis, we try to compare, "Okay. What do you get for your dollar? Who's doing more good per dollar?" And we're trying to look for a way that we can compare the charities we have rather than looking for one universal metric because we think that is a more reliable, robust way to do it. So in the case of how do we compare cash and deworming, what we basically came up with is, if a child is dewormed, there's this hope that they earn more money later in life. That's the main, most impressive effect we've seen on deworming in any of the studies. And so you can kind of look at that and you can say, "Well, that sort of like an investment." You spend a certain amount of money now deworming the child. And then after a certain amount of time, that investment pays off in more money. And then you can look at cash the same way. You can say, "Well, you give someone some money now and then they're going to get a certain rate of return on it plus they get to consume the rest of it."



**0:24:18 HK:** And so we kind of denominated everything in dollars, or we looked at return on investment. Not strictly in dollars because we also adjusted it for this concept of, "Alright, how much is this wealth people are gaining in relation to their income?" And so that was the basic metric we used and by that metric, our best guess is that deworming is doing more good per dollar spent than cash transfers, and that's based on our understanding of how much deworming causes people to make more money, what the right guess for the rate of return on investment for cash is? And then how good the evidence for both is and how much you should be discounting things based on that.

**0:24:55 HK:** With that said, our calculations are very rough. There's a huge amount of room for error. I mean, Elie did a version of the spreadsheet. We basically... We published a spreadsheet where you can plug-in your own assumptions and Elie did a version where he kind of put in two different sets of reasonable assumptions. They got results, but that differed a lot, by like a factor of 20 for SCI and like over a 100 for Deworming.

**0:25:12 EH:** Like 300 for Deworm the World.

**0:25:14 HK:** Yeah, a fact...

**0:25:16 EH:** So you can make what you will of a tool that will leads to results that differ by two orders of magnitude.

**0:25:21 HK:** So I mean, we did this analysis and we're not internally united on how to think of this. I mean to me, it's like if you have Deworming looks about three X's good, but the calculation is really sketchy. I'm like right on the fence about how I want to weigh that, which one I think really does more good, given that I'm more confident in cash and Elie's more of the mentality that we should just ignore that calculation when deciding between the two.

**0:25:43 S?:** You did the same thing for malaria?

**0:25:44 HK:** So, we didn't do that for AMF this year because they are not a recommended charity. Last year, basically we compared bed nets to deworming using a, like a life saved equivalent or a healthy life year saved equivalent metric and found that they were in the same ballpark. So it's like we've got bed nets are a little better than deworming, they're in the same ballpark. Deworming is about three X's good as cash and that seems to be about as much as we can usefully say. I think converting all three into one master metric, I think it's just not going to...

**0:26:14 EH:** What's going on is we're trying to go as far as we can with quantification. And there's some debate about how far it really gets you internally among the staff. And then you're just left with this decision about which charities do you recommend and there's a lot of softer stuff that starts to filter into this equation. How confident are you in the organization's results? How confident are you in the communication you have? What do you think the upside is of giving? And our approach is to quantify what we can. And then for everything that's not quantifiable just write out about it and be as open as we can about the factors that are influencing the decision. Yeah.

**0:26:48 S?:** Are there any interventions that you think are massively funded, but really off based? Or, the money should be really re-directed elsewhere?

**0:26:56 HK:** Sure, so the question is: Are there interventions that are massively funded that we

think are off based? The money should be elsewhere? Yes, although this is an opinion, it's not some... Basically we spend all year trying to find the best giving opportunities, right? And that takes a lot of work. Understanding one charity takes a lot of work. And so it's an organizational strategic decision to not spend time proving that something is bad. It's actually much harder than proving that something is good, 'cause you don't get the same access to information, and it's also, it's not as worthwhile because you can't do as much with that conclusion. Like, at least this way, we can tell people where to give.

**0:27:30 HK:** So it's not a priority of ours. I mean if I were to say interventions that we think are over funded. I mean, I think that there's a lot of interventions where there's just kind of no evidence base. I mean, people have no idea if they work. I think there's stuff like microfinance where I feel like there was a point at which microfinance was an amazing investment and was a great deal. This is like loans to very poor people, but right now there is a lot of money going into it and kind of... There's a lot of assumptions that we just need to like get loans to as many people as we possibly can. And all the really good evidence that's out there is just not supporting the idea that this is making peoples' lives better in a systematic enough way.

**0:28:11 HK:** There is stuff like giving livestock where we've kind of asked, I mean it's not isn't necessarily bad, but we've kind of asked, "Why give someone livestock when you could give them cash?" A lot of times the two are used in the same way, anyway, except the livestock carries a whole host of additional challenges and problems. There's no evidence base there, but more broadly, I mean we also think that interventions targeting the global poor are really good. Your money is likely to go really far and interventions that are not targeting the global poor probably your money's not going nearly as far. So could you be getting something for your money? Yes. Is this a terrible thing to do? No, but I mean that's a big reason for our emphasis on international aid.

**0:28:50 EH:** Yeah.

**0:28:52 S?:** So you brought up this charity Deworm the World, which seems like a sort of an infrastructure bed in some sense, right? And I wonder, in the context of the problem of scaling where it seems like a lot of the problems that people are going to face are in the presence or absence of either human capital or actual infrastructure on the ground, like roads or something. How you guys, if you talk a little bit about how you weigh a charity that is infrastructural in some sense that deals with governments and not directly with governments or building things or doing research or something like that?

**0:29:23 EH:** So are you asking, Deworm the World is playing this like negotiation role as opposed to a delivery role, and how do we think about that question?

**0:29:32 S?:** How do you think about those kinds of indirect charities, in the sense...

**0:29:34 EH:** Yeah, I mean the big trade-offs there, I mean, ultimately it comes down to, how much good do you expect to be accomplished over some period of time based on this charity's activities? And so with Deworm the World, it largely is, do you believe that their activity leads to governments... They work in India, so Indian state governments undertaking deworming programs, running them longer and better than they would have in Deworm the World's absence. And I think this is something we're just starting to think about, because up until this year, GiveWell, at least the traditional side, the international side of GiveWell, had a very clear focus on delivery. It was like,

"We can't, we don't know this stuff. We can't figure it out, so we're going to stick to what we know." And Deworm the World was the first attempt to get at that. And so we're trying to ask questions like, if we talk to people who know the field well, do they say that in general, Indian state governments rely on entities like Deworm the World to make this happen?

**0:30:30 EH:** The obvious trade-off, the reason that this is attractive even if it's much less definite, is that you have the opportunity to get a much larger return, so to speak, on your giving if the organization is successful. Just very simply, SCI spends about a dollar for each child it deworms. Deworm the World spends about three cents for each child it deworms. Now the entire cost of the deworming in India is larger. It's like 35 cents 'cause they're not treating one of the worms, so the whole thing is cheaper. But their money is a very tiny proportion of the overall budget, and so there's some possibility that they're spending significantly less getting this impact, and so you'll get a much larger return. Someone over here? Oh yeah, go ahead.

**0:31:18 S?:** So, I work in IPA and I think we're all really excited about your recommendation for this year, especially given just how the RCT on GiveDirectly and Deworm the World is part of, [0:31:32] [REDACTED]. And I think something that's really cool about GiveWell is that you answer a lot of questions that we can't answer, which is about, what charity should I donate to and I guess a big reason I came is just to figure out a little bit more about how you go from finding intervention that you think is working very well, finding the ideal partner. I know you've talked on that a bit [0:31:55] [REDACTED].

**0:31:57 EH:** Sure. Just for people who really don't know, IPA is an acronym for Innovations for Poverty Action, which is a research group that is associated with Yale University, and they do a lot of the randomized control trials of non-health programs in the developing world. So they were working on, they did a deworming trial, a cash transfer trial, and so they're a big part of the research that we've relied on this report. And so the question is, how do you go from an intervention that has some evidence to a specific organization to work with? Broadly speaking, we have found relatively few organization in the scheme of things, running programs that have the evidence base behind them that the charities we recommend do. We published a list of these each year, there's roughly like 60 of them and even many of those organizations are running many programs across their whole, the whole scope of their activities.

**0:32:50 EH:** They're not just doing deworming or just doing nets, and of course, that raises a question. If you gave to that organization, whether you'd really be supporting the program that have the evidence or supporting something else. So there's, first you're starting with a relatively small universe of possible organizations, and then the biggest thing that, I think, potentially holds organizations up in dealing with us is that we look for a very high degree of transparency and information from the organization. That means them sending us information about what they do, but also them giving us the time that we need to really understand them. And we want to do that because we are, ourselves, still at the stage of learning about what makes a good charity, understanding these challenges of scale, and the only way we can do that is if we have an organization we can learn from.

**0:33:33 EH:** I hope that as GiveWell's influence continues to grow, we'll see what we've seen some of so far, which is organizations having more incentive to want to share information with us because that comes with the possible reward of a significant amount of funding. We also hope that this potentially creates the GiveWell money moved, which is what we call the money that goes to

the charities we recommend. The incentive for people that create organizations that are running the programs that we think have the most evidence or others say have the most evidence behind them. And so we also have a list of those priority programs that we published. Go ahead.

**0:34:10 HK:** Yeah. I mean, I think it's important to be aware of the chicken and egg problem, that GiveWell's based throughout our existence, which is that we need a lot of information from charities in order to make good recommendations, but charities are only going to share their information if they think that our recommendations will drive money. Our recommendations can only drive money if we have them, and if they're good, and if they're based on information, so there's kind of a circle there. And I think we've been working our way out of it and I think we are making progress on this, but the way that the GiveWell dynamic plays out is, we recommend the groups that we can get information from and that we can learn about, and we don't recommend the ones that we don't.

**0:34:45 HK:** So we don't, we kind of don't make wild guesses and say, "Well, this one's probably good." That's not the way that GiveWell operates. GiveWell operates on recommending things that we're able to have a window into and see, and then the money has gone to the ones we've recommended, and then we're, over time, able to say, "Okay, now it's actually worth it to participate in GiveWell's process." We can demonstrate this money moved and therefore there's an incentive to participate. And I think we've seen real progress on that, I mean, just about all the organizations we've worked with, I think you could at least argue that they're sharing more information than they were before than they would, if not for the existence of GiveWell and if not for the money driven. So I think it's led us to some strange phenomena in the past where we kind of have a biased organizations who work with us, but I think that problem is getting less and less severe over time. Organizations are more and more interested in engaging with us, they're sharing more and more information, and so I think the quality of what we know and what we're able to recommend is on an upward path.

**0:35:39 EH:** Yeah.

**0:35:40 S?:** So in the past, you've talked about like you tried to measure by total cost, not by the marginal impact of the money that's donated through GiveWell. Do you do that with regard to government spending as well in the Deworm the World case?

**0:35:53 HK:** Right.

**0:35:53 S?:** And if so, how do you... You were talking about like, where you have the chance to drive much more impact potentially?

**0:35:59 HK:** Yeah. So the question was, how do we... When we look at something like cost per life saved or I am just giving that as an example because we don't actually have life saved at the current top charities, really. When we have something like that, what do we do about that that cost side of the equation. Is that just the money spent by GiveWell donors? Is that just the money spent by the charity? Is that money spent by all parties? And what we've always done, in the past, before this year, is we've counted all costs because we don't want to be vulnerable to gimmicks where... Okay this donor comes in and pays for this part and then someone claims to us that we're only paying for a fraction of it and stuff like that. And that's something that we've written about.

**0:36:37 HK:** But this year, I think, we tried with Deworm the World, we knew that this leverage case is such a big part of their case. And you really have to believe it to believe they're good and it really is quite plausible. So we really tried to get at it this year and we ended up taking our best guess at a leverage factor. So basically, Deworm the World, spent something like \$1 for every \$10 of total expenses, in the program. So do you want to think of that as 10 to 1 leverage? And we think that would be going too far. We talked to people who'd been involved in the program, not just from Deworm the World but from the Indian governments. We asked them, these are very imperfect questions to ask someone, but "What would things look like without Deworm the World? What's going to happen if Deworm the World disengages? And things like that. And we have written up the results in our site visit notes, in our review, and now we're willing to take a guess that some where between 1X and 10X for leverage. We do that in our cost effectiveness spreadsheet and I think most of us are somewhere around 3X.

**0:37:34 EH:** Yeah, I mean we really don't know there's like a... In our review, there's really like a footnote with different staff members saying what they think. I mean we have no... We honestly, don't know and we're just very tentative [0:37:42] [REDACTED]. So I think we should pause and just have the second half, which is about the newer areas we're going into. And then after that, we can take questions on that or anything else that folks are interested in.

**0:37:52 HK:** Cool. So I'm going to give a little bit of a summary of what we've done with GiveWell Labs this year, starting with what GiveWell Labs is. So, basically, as you probably can pick up from the conversation we've been having, what GiveWell has historically focused on is these very direct interventions, these charities that are kind of what one might call "giving is consumption", where your sort purchasing lives improved or lives saved, and what we've been doing for the last couple years, but really intensifying this year, is thinking about, "What would we do if we started over?" And we said, "Let's think, instead of from the perspective of someone giving, let's say, a \$1000 and wanting that money to pay for a certain amount of bed nets, what if we think from the perspective of more of a major donor who wants to take big risks?" Possibly create organizations, possibly engage in political advocacy, do a lot of things that are less measurable, less direct, less linear maybe higher impact.

**0:38:47 HK:** What would we do if we were thinking that way? Is there a way to apply our framework and think about how do you do as much good as possible with your money and not just limiting yourself to things that are proven and cost effective and scalable? That's called GiveWell Labs. It's called GiveWell Labs just to denote it as an experiment, and this project really came about... It was definitely inspired by our connecting with Good Ventures, which is represented here by Cari, and which is going to be a major foundation, thinking about this question, struggling with this question and we've been working very closely together on this question, basically working as a team to answer this question.

**0:39:25 HK:** Our motivation, for doing this? First we think that we've spent a while trying to understand these sort of linear giving opportunities and we feel ready to explore other things and we feel we may be able to do better, and also, we think that we're going to be able to influence larger donors more successfully if we're looking more broadly at all the things that a major donor can do. This is not only for larger donors. This is something where if we find good giving opportunities we'll also think about how is it that a smaller donor could participate in these, perhaps by pooling their funds, perhaps we'll find things that are a natural fit for a smaller donor. But we're definitely going at it with this mentality of there's a lot of money to be spent how do we spend it as well as

possible?

**0:40:04 HK:** So that's what GiveWell Labs is. I have spent most of my time on it this year, Alexander has spent most of his time in it this year. And a major goal, of this year, was not only to build capacity, hire new people for our traditional work, but also to spend a lot more time on GiveWell Labs that we have in the past. So, in terms of where we are and what's happened there's a fairly recent blogpost, it's called GiveWell Labs Update. I think there's probably more than one post called that, but it's the most recent one.

[laughter]

**0:40:33 HK:** And it is a pretty good summary of where we are but I'm going to summarize it so that... One of the biggest things that we've kind of changed from the beginning about how we think about GiveWell Labs, is that we sort of changed our unit of analysis from the project or the organization for the cause. And so I think very early on, when we were engaging in this work, we would try to look, "Where's a good organization. Where's a good project we could fund? And we noticed a lot of problems with taking this approach. One of 'em is that it's hard to evaluate a project without learning a lot about the space and the context within which it sits, so if someone brings you a proposal for a new kind of malaria drug, it's like well what are the other malaria drugs being worked on and what's going to be the use of this? There's a lot of other things you need to know and once you know them there's a lot of other overlapping projects you're able to look at.

**0:41:23 HK:** So, it felt like kind of a... We saw a lot of projects that we couldn't evaluate. It just felt like it was a less clumsy way to evaluate than was to going to be to think about causes. Another reason that we've found this helpful is that people often will not necessarily have project proposals lying around unless they think they can get them funded, and so people will often sort of, they will start thinking about something they could do or they will start working on telling you what they could do once they believe they have a high chance of getting funded for it. And so, this makes it very difficult to take an approach of saying, "Well let's look at everyone's ideas and choose the best ones," because a lot of times the formation and the writing up of an idea depends on what people think they can get funded.

**0:42:07 HK:** And so, with that in mind, we've kind of switched our approach and I feel that it's really paid off kind of trying to look at causes. And I define a cause as basically a unit of intellectual specialization. So it's like some kind of problem with the world that you think you might be able to do something about, such that you can study up on this problem, get to know the people who think about it and you get sort of returns to scale, such that once you get up on the cause, you're able to evaluate a lot of the people and organizations and causes in the space efficiently. So, a cause could be something very broad, it could be something like climate change or it could be something narrow like a sub-cause, like something like geo-engineering which is an approach to climate change that I'll talk about in a second. So, it isn't necessarily broad, it isn't necessarily narrow, but what it is, is it's kind of a problem one could work on as a philanthropist and a way of describing what one is interested in.

**0:43:02 HK:** So, now that we have that frame, the basic approach we're taking, we have a list of something at this point like a 150 potential causes, and these range... Everything from kind of very specific ideas about policies that ought to change or technologies that ought to be developed, to very broad areas where we just... Malaria control. We know this is important, we're not sure exactly

what to do. And we have this kind of process where we have what we call a shallow investigation of a cause. Actually, let me take a step back. So, what we're trying to do with the cause is we're trying to ask three key questions. One of them is, what is the problem that people are trying to address? Two is, what can be done about it? And three is, who else is working on this? And we're looking for causes that are important, but underfunded relative to how important they are. So kind of the ideal for us is if you have two causes, they seem about equally important, but one of them has 10 times as much philanthropic money and philanthropic involvement, you want to take the one that's less crowded and that's been one of the heuristics we've been thinking about using.

**0:44:05 HK:** So, with that in mind, we basically do shallow investigations, which mean talking generally to one to three people and just getting a very basic sense. Okay, what are people doing in this cause? Who works on it? How crowded is it? And then, we've done something like 20 or 25 of those, about half of which are public and about half of which are like 90% done and we're just waiting on someone to sign off on making something public before we publish them. Then, we do medium investigations, which is, we take of the shallow ones, we take the ones that look more promising and we'll talk to instead of one to three people maybe something more like 10 to 30 people and really try to get a landscape, where we'll try to get a more comprehensive picture of who works in this space, who is doing what. And then, from those, we take the ones that we think are most promising and we're experimenting with this kind of what we call a semi-deep dive, and these categories are very much inflexible, we change them all the time because we're experimenting. But it kind of consists of being ready to commit some funding and so far there've been a couple of grants that Good Ventures has committed to based on this idea of being exploratory and of making some grants that you hope to learn by following.

**0:45:18 HK:** And also, just the mere act of saying we're ready to make some grants dramatically changes what ideas emerge and how people react to you, so, being ready to do that is very crucial. And then finally, there's a level of investigation we haven't done, which is really committing to a cause and saying, "Okay, we're in this for a substantial amount of time and a substantial amount of money." I think you would get further interesting results if you did that, but I think we're not ready to do that with any particular cause yet. We're not doing these processes one at a time, we're doing them in parallel, and that's very important to understand how we're thinking because in some theoretical sense, it might be smart to do a shallow investigation of all 150 causes that we've ever heard of, then pick the ones through the medium, then pick the ones through the deep. But I think, in some ways that would be easier to explain and more systematic, but I think that would get us much worse results because what we learn from each stage of the process feeds in to how we do all the other stages.

**0:46:16 HK:** And so, we think it's really important that we move in parallel and for that reason, a lot of what we're doing is kind of, we pick things based on whether they seem promising to us, based on whether we have a good plan for attacking them, whether we have staff who are excited to go after them, whether we think that investigating them will teach us a lot of things that we can use for other causes. And so, it's kind of a mishmash and it's not terribly systematic at this point, but we expect it to get more systematic over time. So right now, it's like we have a list of 150 causes, we'd pick the ones that looked best by those criteria, we look at them, we picked the ones that looked best after the shallow. It's not like we have a formula and I think at this early stage of the process, a formula would be a bad thing and it would actively get in the way, but we do try to write up how we're thinking.

**0:47:01 HK:** So, with that basic process in mind, I'm going to quickly run through some of the causes that we are at the latest stage with, these are causes that we're kind of potentially interested in making some small grants, so these are ones that have really stood out to us. Labor Mobility, so one of the arguments we've heard multiple times, and I think is somewhat supported by evidence, is that the most dramatic and effective and proven anti-poverty intervention that exists is helping someone move from a poor country to a rich country. That their labor actually becomes worth far more, like 10X as much, they'd make a lot more, they often send the money home. This is an excellent anti-poverty intervention and it's something that is not generally seen as a humanitarian issue. And so, obviously, there's a lot of debates about immigration, about immigrant's rights, that's all over the news right now. We are interested in the particular angle of this that has to do with helping people go from poor countries to rich countries as a humanitarian intervention, and that's been something that's interesting to us based on arguments about the importance of the issue more than about the tractability. We think it's a really hard one.

**0:48:09 HK:** Another issue that is a very different profile, very different reason we're interested in it, is Criminal Justice Reform. So this is the idea that basically we have incredibly high incarceration rates in the United States. It doesn't seem like those incarceration rates are really necessary to have the low crime levels that we have. This is not an issue that we think is the most important issue in the world, but it's an issue that we've been pointed to repeatedly as something that's moving, as something where there's... One of the very few issues where there's kind of the beginnings of a potential by part as inner trans-partisan consensus, where things may actually happen especially at the state level in the short term. And so we think of this as an excellent learning opportunity and we're looking into that a bit.

**0:48:53 HK:** Another cause, geo-engineering, this is kind of on the... Along the lines of climate change, we looked into climate change a fair amount, we think it's a really important issue, but most aspects of climate change, the things you could do about them, there's a lot of philanthropy and a lot of people are working really hard to do something. And one of the areas where there doesn't seem to be a lot of attention is this idea of the worst case scenario, what we're going to do about it and whether we could develop technology that basically is very dangerous and risky in in of itself, but could bail us out if the climate change projections are completely off and if everything's far worse than anyone's predicting.

**0:49:30 HK:** So it's kind of a worse case scenario technology for climate change. The idea we're investigating is not to start deploying these technologies and shooting them into the air but rather to increase the amount of research and analysis that's going into them. Because if that day ever comes and this technology starts to look like something that people want to do, the more informed about it we are the better and the fact that it's a really tiny area of research right now does not seem like a good thing. Then other causes, malaria control is an example of something that is, it's got a lot of money, there's a lot of people paying attention to it. It's not an empty space at all and yet bed nets seem like a really good buy. And so we want to see what else might be a really good buy and what might be worth doing in malaria control.

**0:50:11 HK:** Finally we are pretty interested in this issue of factory farming and the unethical treatment of animals. There is no consensus on staff about how to weigh animal suffering versus human suffering. That is completely not something that we agree on or that we're sure about. But that isn't the question we're asking right now. What we were noticing is that according to a certain value system, this is an incredibly important issue and although there's lot of non-profit dollars on



animal welfare writ broadly, there's not a lot of attention. Most of those dollars are not dealing with factory farms.

**0:50:43 HK:** And so what hopefully you hear in this list of causes is a variety of different things that make a cause standout. We don't believe in having one criterion, "Oh We work on the most important problem or we work on the easiest fix." It's like some of these are on there because they're empty and important, some of these are on there because they're empty and tractable, some of them are on there because they're crowded but they seem really good anyway. And so it's like we're looking for things that stand out by different criteria and exploring them in the hopes that we'll learn more doing that.

**0:51:10 S?:** [0:51:10]       .

**0:51:12 HK:** I think I gave enough. I think we just got to move on. There's more stuff we're thinking about, a lot more yeah.

**0:51:13 EH:** Are you done now?

[laughter]

**0:51:24 S?:** I'm so excited I kind of want you to read them all.

[laughter]

**0:51:28 HK:** I mean if you want. Let's take some questions.

**0:51:31 EH:** We have a website.

**0:51:32 S?:** Yeah.

**0:51:33 EH:** So, questions.

**0:51:33 HK:** I mean the blogpost GiveWell Labs Update is a good place to look. Yeah.

**0:51:38 EH:** Go ahead.

**0:51:43 S?:** So we got a question that can apply to actually both halves of this discussion. So one of the things that I'm clearly aware of as a donor, I know you guys are even more aware of this, you are trying to find out information of just how hard it is to get information that's really accurate about the impacts and to predict what the effects of any given program are going to be. And one thing that I think that is really important especially when doing interventions that serve communities that maybe hard to communicate with which may be overseas or anything is that people that are directly affected by the intervention have or probably a lot to say about like what the intervention is doing for them, which may not be something not easy to predict as like an outsider, looking at that.

**0:52:33 HK:** Right.

**0:52:34 S?:** And I'm curious to hear more both for the narrow causes that you have identified in the

first half and also for like moving forward in the GiveWell Labs case, what matrix do you use to judge charities or causes, based on this metric on how communicative they are with the affected communities and what are your standards around that and how do your charities perform?

**0:52:59 HK:** Right. So the question is kind of how are we integrating into our research this idea of the constituents and the beneficiaries of charitable programs having a voice and saying how they feel about what's being done. I mean, I think I would say broadly that where this kind of analysis can be done and where this kind of information can be collected, it should be and it should be considered extremely important. At the same time, I think that we have seen a lot of limitations to the practical ability to use that kind of information to get very far. A lot of times when people get free stuff they want to say good things about the program, a lot of times when things are not structured that way they're less likely to want to say good things about the program. And so I think this is not... I mean this is not frankly something that has come up a lot in the charities we've have looked although like I think it is something that we have had an eye out for.

**0:53:54 HK:** But good examples of this, I mean GiveDirectly, they give out cash. So they have asked people a bunch of question, about, "Would you rather have the cash in a lump sum or would you rather have it over time? Would you rather have the cash or would you rather have a cellphone for part of the cash?" I think that's actually... It's a little bit hard to be that intervention in my opinion if what you are most concerned about is kind of people's attitudes towards the help they are receiving. Because I think they really are tailoring it, according to what people are telling them they want, and what people mostly want, is they want the cash, they want it in a lump sum. And I think that makes sense, so there is that. Then it's like how do you trade that off against deworming where kind of the whole idea of the intervention is it's incredibly cheap to do it at a population level. Many of the people getting it don't really understand what it's doing, and you are only able to pick up effects using a large sample size and using kind of a complex study.

**0:54:43 HK:** And so, that's a question that I ask myself when I'm deciding which one to give to, which should I believe is kind of... This equation we have, that says it does more good or the kind of a obvious, people like cash, and I think that's an example of where the question is tough to answer. I think it's also tough in some of the causes I just gave for example, I mean its hard to apply this kind of analysis to planning for climate change. But it is something that I agree is very important. There.

**0:55:09 S?:** Criminal Justice Reform, I'm obviously very much on board with funding around. But it seems very US specific, and I'm wondering whether you're getting any hooks into issues that other... It seems like there might be, if you can get financial reform in Norway or something, I'm like making that up. But it seems unlikely to be the Criminal Justice Reform in the US, is like the best of such intervention globally.

**0:55:33 HK:** Yeah. So, the question is, are looking at basically these policy issue like Criminal Justice Reform, are we looking at the US or other countries too? I think in theory, like I'm totally with you that we want, like in theory you'd want to look everywhere. And I think for some things like Labor Mobility, that may even be something we do in the short-term. But I think the other thing is that we're optimizing this agenda not just to find the best stuff quickest, but also to use our time well and to get good at understanding this stuff quickly. And frankly, we are just much better positioned to understand the US and will move a lot faster, and will learn a lot more about issues that I think will then service well if and when we go into other countries.

**0:56:14 HK:** And that's been the consistent feedback we've gotten when we speak to people who know more about political advocacy than we do, it's like you guys have to get up to speed on how this works, you can't spend your first step going off and learning another culture in another language, and I do agree with that. So this is a pragmatic decision to focus on things that... It's much like the original decision when we started GiveWell, to look at this proving cost effective stuff. It's like always trying to bite off something that we think we'll actually be able to chew, always trying to be very ambitious with that. But still believe that we are going to be able to chew it before we move to the next thing. Yep.

**0:56:50 S?:** In doing your multiple levels of work of study for GiveWell Labs. How often have you been surprised when moving from shallow study to maybe a medium depth study about the effectiveness of intervention in that area, either on the up side or the down side?

**0:57:03 HK:** Sure. So, the question is how often have we been surprised moving from one level of investigation to another within GiveWell Labs? About how good an intervention looks or how good a cause looks? We've had discussions about this, we made little lists of like what we expected before and what we know now, and that's something we'd even probably write up if someone bugged us to do so. I would say in some cases we just haven't done enough to really be surprised, but in some cases we have been surprised. So, examples would be, I think Criminal Justice Reform is just something that was not on our radar when we started off. And has kind of looked better as we've gone more into it, and we've just seen that the opportunities to get something done just look qualitatively different from most areas of politics.

**0:57:46 HK:** Conversely something that was very high in our list a while ago was kind of this meta research idea or this open science idea of data sharing and open post-publication peer review. And I actually I don't want to go over board here, 'cause I still think that's a really interesting cause and one that has a lot of potential. And there are things that are related to it that I think have a great deal of potential, but it's not... I think that the cause was a little bit more crowded than we thought it would be, and it isn't as high in our list now as it was before, I don't want to go too far with that we may end up in it anyway.

**0:58:24 EH:** I think given this stage we're at, a lot of what we've learned are going from knowing nothing to have completed the first level of investigation. So, to give two quick examples, we briefly looked into the issue of family planning, basically helping people have access to contraception in developing countries, and issue seems important, we should try to do something about it if we can. But there is a great deal of philanthropic funding already there. And if you went down the list of the major foundations, virtually all of them has a relatively substantial family planning program. That kind of the level at which we know that area, and we still might find things that they missed, but it certainly is an adjustment from what we knew before we had looked into it at all, which is this is just a large problem similar to many other global health problems. And really unique among global health problems, this issue has a lot of philanthropic funding in it. On the flip side, something like developing world infrastructure, roads, power, things like that, there is a great deal of money, there is virtually no philanthropic funding. And if you look at, you talk to people on this issue, they often talk about the ways in which the World Bank or the African Development Bank or US regulation cause funds to be used suboptimally.

**0:59:31 EH:** Now, all of this is based on like 10 hours of research, so, it's very little. But it's

something where a relatively small amount of research moves us a little bit in the direction of knowing more about how promising an area looks.

**0:59:47 HK:** Factory Farming was another surprise for us, although probably not for a lot of people. And that probably illustrates kind of how this process goes, which is that, a lot of our shallow investigations get us to the point that like a person who is like ever thought about this would already be.

[laughter]

**1:00:01 HK:** So, it could kind of be a little silly for people who know about an issue to read our shallow investigation and be like, "Yeah," but I mean there's a lot of issues, and so to know that much about a lot of issues is what we're after. So, Factory Farming, I think my casual belief was like, man, animal charities are enormous and there's a billion of them, but it turns out, it looks like, that they're actually are mostly focused on other things.

**1:00:25 EH:** Yeah.

**1:00:27 S?:** Can you explain a little bit more about as GiveWell goes forward, how are you going to balance resources between the original work that you talked about in the first half and then GiveWell Labs, and I guess part of the reasoning behind my question is that you're doing kind of pioneering work in meeting the original goal of finding the best most effective charities to recommend, and I think based on what you talked about earlier, you're still finding... You're still an experimental process, you're still learning a lot about how to do that work most effectively. So, I'm just curious to think, how are you thinking about the mission of GiveWell more broadly, and how you're going to allocate resources across these two different areas?

**1:01:10 EH:** Right. So the question is about the future of GiveWell and balancing resources across labs versus the more traditional work. I mean this is a tough question we have been struggling with over the last year and I think we will struggle with in to the future. The quick answer is that, I think, GiveWell Labs is more likely to be the bulk of what we focus on in the future, we want the work we've done so far to exist, and improve, and be a great resource, and I think it will continue to be that. But if I had to guess like where will the majority of the effort and funds go to 10 years from now, I would bet on substantially... So like labs being larger than the GiveWell traditional stuff.

**1:01:54 HK:** I might put it a little differently. I think that's probably true, but I think another thing worth saying is that it's also a matter... We're trying to take the traditional work and turn it into something that's more self sustaining because when we were originally doing the traditional work, we were just completely making things up as we went along. And now, we have like sample charity reviews and interventional reports, and there's kind of a process and there's kind of some procedures that still... There's a lot of judgment in it, but it's something that we can envision making it the kind of thing that we have that is kind of running without needing Elie and me, the co-founders of GiveWell, to be doing all the work. GiveWell Labs now looks more like... It's where I spend the bulk of my time, and it's kind of where the tradition of GiveWell used to be. It's completely improvisational, and I think it's going to take a longer time to get to the point where we can kind of have it be something that doesn't need as much involvement from me, for example, to run.

**1:02:58 S?:** Yeah?

**1:02:59 S?:** So, can you kind of like... I know this is going to be... Its case-by-case basis but can you give us like an example of how you, given that you pick a topic, how do you go about investigating it? How do you do it? Is it like you spend 10 days, when do you like say, "Okay, I've had enough?"

[laughter]

**1:03:20 HK:** Right.

**1:03:20 S?:** What's your... A day like? It's just a logistic question.

**1:03:26 HK:** So the question is how do we logistically investigate a topic? I mean, obviously, the answer is that depends. It's very different answers for the two because for traditional GiveWell, it's like we've certain processes that we've come to believe are very important and I could spell them out for you and tell you how long they take. And for GiveWell Labs, it's like changing faster and it's more... One very high-level general principle we follow is that you don't want to lock yourself into a process before you're pretty confident in the process, and so a lot of what we've done with GiveWell Labs, it's like the less we know about an area, the more we're willing to kind of just walk around, immerse ourselves in it. Scientific research is one of the harder things for us to look at and so I've done things like reading part of a biology textbook and just like hanging out with scientists, and like talking to them in kind of undirected ways until you start to feel like you understand what the key questions are, you start to be able to impose a framework.

**1:04:23 HK:** So that's a general principle, it's like the less we know about something, the more unstructured we are and the more we kind of want to look around especially talk to people who know the issue but also know us and are going to be willing to spend a lot of time talking to us instead of setting up a bunch of high-stakes meetings where we have to make the most of an hour. So, that's a general principle and specifically, what we're doing right now, I mean a shallow investigation is one to three, basically to simplify, as one to three conversations aiming to answer the questions, what is the problem, what can we do about it, and who works on this. And you're trying to choose the conversations for maximum efficiency. So, we actually have like a whole list of like this is what makes a good person to talk to. It's got to be someone who's going to know the answers to these questions. Here are some predictors of who's going to know the answers. Then a medium depth investigation is like exactly the same thing, but like roughly 10 times as many conversations. And then, a semi-deep dive is basically that plus being willing to make some moderate-sized grants.

**1:05:27 EH:** Yeah, in the back?

**1:05:29 S?:** So, I know you guys have done some work looking at philanthropy historically to see what's worked. I was wondering how much, if at all, what you learned there has guided your decisions about what to look into and like your [1:05:40]            does?

**1:05:42 HK:** Sorry I just...

**1:05:43 EH:** Yes. Oh, god. So, the question is about... We have this project we call the History of Philanthropy and so the question is how much is that informed of what we've done so far. So, the

idea of this History of Philanthropy project is there's just surprisingly little information about what philanthropy's track record has been historically, what it's achieved, what it's failed at. There is some information about this, but very little of that would be, I'd say up to the standard that's really convincing to us that this philanthropist or activity really had the impact that it purported to have had. We've really just gotten started on this. It's one of the areas where most... I am most excited about with GiveWell Labs because I think it has the potential to influence what we do, what others do down the line with this type of information. Thus far, I think most of what we've done is just like start to figure out the process of how we'll gather this information understanding better what's actually out there more than we've used it to inform the decisions that we've made.

**1:06:44 HK:** Though I would say that claims about History of Philanthropy have influenced me a great deal. And it's more that we haven't really vetted those claims enough for that vetting to have made a big difference yet. So, I actually I'm much more positive on the idea that philanthropy has had a lot of impact than it was before I started try and answer this question. I've asked a lot of people, "What's been accomplished?" I used to ask this of like people who work at foundations today, "What has your foundation done?" And that did not give me very good answers and I just didn't really think, like I didn't know that there was much. But asking this question a little bit more broadly if a broader set of people looking for books on it. We have a post called Philanthropy Success Stories or something like that. And just very informally, I have come to believe that philanthropy has had a lot of cases where someone's kind of decided they want an idea or a cause to get more attention. And they've successfully put it on the map, and that has successfully made a difference to how governments behave, how other funders behave. I could name a lot of cases where I believe that to be true. I can't name any yet where we've like gone all the way in and say, "Yeah this is definitely how it worked".

**1:07:51 EH:** I think broadly it's interesting because a lot of the claimed successes are these big home runs. It's like, "Small investment, giant impact". Right? And then we're thinking about our GiveWell top charities and we say, even in these cases where you really are doing almost the simplest thing, there's questions about impact and I think that certainly influenced me. I think it's influenced our thinking about the relative, the trade-offs between "high risk, high reward" versus "lower risk, lower reward" charitable giving. Yeah.

**1:08:23 S?:** Some of the steps we're doing at GiveWell Labs sounds like the kinds of things that other big foundations would have done like the Gates Foundation [1:08:29]            might do for various causes or something. How do you know that what you're doing is novel and then like have you talked to other big foundations about whether they've done things like this?

**1:08:38 HK:** Sure. So, the question is how do we know... A lot of what we're describing could have been stuff that the Gates Foundation did. How do we know they didn't? How do we know we're doing is new? So we have spent a lot of time talking to other foundations, that includes the Gates Foundation, that includes other foundations. Cari has spent particularly a lot of time on that so if you want to get her thoughts afterward. And so we have done a lot of that and I would say that my impressions at this point are that there is going to be a fair amount of common ground between what we end up doing in what other foundations have done, I think that makes sense. I don't think that every single thing foundations do is just crazy. I think most of it has reasons behind it and comes from somewhere I think there will be differences too. Two things that I think are very different, one is I think looking this broadly at this many causes and saying which one's are the best, I really think is, if not unique, then very unusual and possibly unique.

**1:09:35 HK:** And that's something that I having talked to a bunch of foundations, I really do believe that. I think the more common approach is to say, "I'm passionate about X and just go right in." And in some ways that's a lot easier but this broader approach I think it is at least highly unusual. And then the part that I know is unusual and know is unique is the part about talking about it in public because if someone had done this before and written about it then we wouldn't even have this question. And so, hopefully the next person who wants to come along and do this won't have this question. And we'll be able to get up to speed a lot faster, doesn't mean they'll do what we say, doesn't mean they'll agree with us on everything. But there's basic questions we're asking that other people shouldn't have to ask in the future.

**1:10:16 EH:** We're going to do two more questions. So here and here. We'll stick around after to ask any... Answer more if people want.

**1:10:24 S?:** Sure. I also want to know like what kinds of things that you guys are doing to grow the amount of money that goes towards the charities that are recommended by GiveWell?

**1:10:35 HK:** So, the question is, what are we doing to grow the amount of money that goes to the charities we recommend. So, this is a metric we call "money moved". It's one of the key ways we assess our impact and we read about this in our annual evaluation. Historically we've found that we've been very bad at proactively increasing the amount of money that goes to our top charities via marketing or something like that. And instead the factor that has always driven more donations, more media attention, what have you, is doing research that is high quality. Having people recognize that and then them if they're in the media writing about it, if their donors telling their friend's about it. That's always been the more successful path. The best example of this is the philosopher Peter Singer who wrote a book called "The Life You Can Save" which talks about giving to charity and he points to us he is still constantly in the media talking about charity and he points to GiveWell, and that's been better than any marketing department that we could've had.

**1:11:36 EH:** At the same time we also have seen our growth be pretty good like we've roughly doubled the amount of money that's gone to our top charities every year that we've been around and I don't know, if growth will quite hit that this year, but it's still been reasonably strong. And so at the same time we see, not San Francisco start-up growth, but pretty strong growth rates that we're happy with and so we want to allocate the time to research over marketing.

**1:12:03 S?:** So this is kind of a methodology question, but it seems like the historical portfolio of investments looking at it from a dimension of time that needs more lower risk and return is a pretty near term, I mean around deworming, very quick. And then you take, at the opposite extreme, like geo-engineering. How do you guys think about the time against the kind of "teach a man to fish" question, education versus direct cash assistance, geo-engineering versus energy efficiency retrofit would be the equivalent. Do you apply discount rates? How do you think [1:12:37]        about time?

**1:12:37 HK:** So, the question is how do we think about time, how do we weigh very long-time horizon on uncertain outcomes versus short-time horizons uncertain. I mean I think conceptually the right way to think about this is that things that are more uncertain and farther away get discounted in various ways, and then you compare them post discount. I think in practice to actually write down those formulas is not going to help you very much, and I think we are unusually experienced

at trying to write down those formulas, and see where they get us. So, I think we can say that with some background there, but I mean that is how we think of it. And so, as I think you can see, I mean we're not picking one end or the other, and that's one thing that does make GiveWell and our approach pretty unique I think, is that I think most people decide where they want to be on most of these questions. So, do I want to be short term, long term? Do I want to fight winnable battles or important battles, and we're kind of like...

**1:13:28 HK:** Well, we have one cause that's a really important battle, we don't know if it's winnable. One cause that looks like an unusual winnable battle, we're not sure if it's the most important. Another thing that's completely different and... Yeah, I mean I think... What we're doing right now looks a lot more like being interested in anything that stands out on a particular dimension than it does like trying to get the exact right answer, and I think that approach is going to be better for our long run sort of learning and breadth of understanding.

**1:13:56 EH:** So I think we're going to wrap up. We really appreciate everyone coming and engaging with us in this discussion. We know that this is not quite as swanky as your normal charity event. And we appreciate your taking the time to talk some charity with us and grill us with your questions. It's really part of what we're trying to do. So, thanks. We're happy to stick around as long as folks want to kind of answer more questions one on one, so thank you.

**1:14:19 HK:** At some point, we need to go home.

[laughter]

[applause]