

**GiveWell NYC Research Event November 11, 2019 – Top Charities**

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

**00:00 Catherine Hollander:** GiveWell is a recommender of great opportunities, so unlike many other charity evaluators, we're not a database. We focus instead on conducting independent research to find and recommend opportunities that are exceptional, that can do a lot of good per dollar donated. We don't just focus on identifying charities' public information, but we spend hundreds of hours of research and vetting into the opportunities that end up on our list. We're not trying to be comprehensive. We're really trying to focus on a few excellent organizations that we can recommend.

**00:36 CH:** Because of our focus on charitable dollars going as far as possible, we have ended up recommending charities exclusively that work in the poorest parts of the world. And the reason for that is that we found that your dollar can go the furthest in parts of the world outside of the United States where people are living in extreme poverty. We've also focused on global health and poverty alleviation because we found that the evidence of effectiveness of the programs in that area tends to be quite strong. The impact per dollar donated tends to be quite high. And we've also found that these are tractable issues to work on, so a donation can actually have an impact in working toward improving the world and solving the problem.

**01:20 CH:** We recommend three top charities that implement programs that we believe save lives. These are the Against Malaria Foundation and Malaria Consortium, which we recommend for their programs to prevent malaria, and Helen Keller International, which distributes vitamin A supplements, and we think that can lead to reduced child mortality from infectious disease. We also recommend five charities that implement programs that we believe can lead to higher incomes or higher abilities to buy things, higher consumption levels. So these are four charities that implement programs to treat intestinal parasites, which is a health intervention that we think has income benefits. I'm happy to talk more about that evidence-base there, and GiveDirectly, an organization that distributes unconditional cash transfers to very poor households.

**02:08 CH:** So now, moving into our 2019 recommendation, all of those top charities remain on our list and our top recommendation for donors is to give to grants to recommended charities at GiveWell's discretion, which is an option on our website that you can choose if you're donating to GiveWell where we will allocate it to one of the eight top charities on our list that we believe can use the funding most effectively at the time that we make those grants.

**02:37 CH:** Each of our top charities has been on our list in the past. If you've been coming to these events a while, you might see some familiar names. But once we add a top charity to our list, that's not sort of the end of the road, we don't say "Great, okay, we assume they're wonderful forever. And that's the end of our review process." We actually spend a lot of time each year following up on all of our top charities to confirm that they are still doing work that we consider to be highly effective and highly cost-effective. So at the beginning of each year we sit down as a research team and we ask ourselves, "What is the information that we could look for this year that is most likely to update our recommendations, to cause us to change the way that we prioritize funding among our top charities or even to cause us to remove a top charity from our list or add a new top charity?"

**03:21 CH:** And so this year we looked for information from our top charities such as detailed information on how they are implementing their programs, monitoring information to show that they are successfully reaching the beneficiaries that they're intending to reach. We also look for costs, the cost that they and their partners have incurred implementing their programs and how many people they reach. And finally, we ask them how they would spend additional funding, which is a really important part of our recommendations. We want to know that every additional dollar can do a lot of good. And so simply looking at a charity's track record isn't sufficient for us. We want to know, "Will additional dollars do a lot of good?", so we ask them how they'll spend future funding.

**04:03 CH:** We also speak with outside experts who can inform our views of the programs that our top charities implement. So this year we spoke with experts, for example, that work on measuring vitamin A deficiency rates in the areas that one of our top charities, Helen Keller International, works. Helen Keller International, as I mentioned earlier, is a charity that distributes vitamin A supplements, so the level of vitamin A deficiency in the areas where Helen Keller International works, can have a big impact on how cost-effective we think their work is.

**04:35 CH:** We also spoke with experts in the malaria space to get a better sense of the total global funding needs for malaria nets and a program that we recommend called Seasonal Malaria Chemoprevention, which is the distribution of preventive antimalarial drugs. So we wanted to get a sense of the total global needs for those programs so that we could better understand the impact that our charities that work on those programs are having. We also generally conduct site visits to our top charities. We go and check out their work on the ground. This year we visited Malaria Consortium. This is a picture of GiveWell staff, including GiveWell staff member Olivia Larsen who's here tonight, on a site visit to Malaria Consortium this August. So we go and we see their work on the ground, and this is another helpful time for us to check in and make sure that we're understanding their model correctly and thinking through any potential differences in how they're actually implementing the program to what our understanding was.

**05:34 CH:** We don't believe that each of our top charities is equally cost-effective. We believe that it's important to prioritize among our top charities. We expect that our top charities have total funding needs that exceed the amount that they're going to receive from donors, and so given the number of donations that we have to allocate, we want to make sure that we're starting with the highest priority and sort of moving down the list. So our best guess right now of the highest priority funding gap among all eight of our top charities is at the Malaria Consortium. We think that this is a very cost-effective gap. We think that this is a very strong organization, and we think that they can likely use an additional \$34 million over three years very effectively. So we're planning to tell donors who prefer to give directly to a specific charity that Malaria Consortium is our best guess of where they should give to maximize their impact.

**06:31 CH:** But getting back to prioritization, as I mentioned, our top recommendation is this option called "grants to recommended charities at GiveWell's discretion", because we do spend a lot of time thinking about where the highest value need is among our top charities. We find it incredibly helpful to be able to take this flexible funding and grant it to the top charity or top charities that need it most at the time that it's granted. So, funds that are given to us with this designation will be given to one of the eight top charities on the list. And it really is a helpful way for us to maximize the impact of donors' funds, since we think about this question in an ongoing way. And just as an example of how we've spent discretionary grants, so this option that I mentioned, in the past...

They've gone most recently to Against Malaria Foundation, Malaria Consortium's Seasonal Malaria Chemoprevention program and SCI Foundation, formerly known as just Schistosomiasis Control Initiative. So we update these recommendations every quarter when we grant these funds out. I'm going to take a pause here for questions related to our top charities before I go into the next topic which is some new research that we've done. Does anyone have any questions on our 2019 recommendations or review process?

**07:46 Speaker 1:** Yeah, you've talked a little bit about how you might choose someone within the malaria space. But how are you choosing... How are you choosing malaria, how are you choosing vitamin A? How are you choosing the higher level... Making the higher level choices. It sounds like you need some real health experts to be doing that for you. Yeah so I'm curious.

**08:10 CH:** Yup, so the question... I'm just going to repeat the questions for the sake of the recording and the microphones. The question is, "How do we decide which programs to look for charities implementing?" So we have a few charities that work on malaria, we have a charity that works on vitamin A. How do we decide that those areas are the places that we want to recommend top charities? So the first step of our research process is looking into the program areas that we think might be most evidence-backed and cost-effective. So this is sort of independent of any charities before we're looking at any specific organizations. We say, "What are the areas in which we think that dollars can go the furthest to save or improve lives?" And at the very highest level, I mentioned this sort of very briefly in my overview of what GiveWell is, but we found that those tend to be opportunities in the global health and development space, because we found that there tends to be a lot of independent academic research into the interventions in that space, and also that they tend to be some of the most cost-effective ways that we're aware of to help people.

**09:07 CH:** And so in an ongoing way we keep an eye on new studies that are coming out. We watch the academic literature, we speak to experts who work in global health and poverty alleviation and look for things that seem promising, and once we identify something that seems promising we operate as sort of a funnel where we spend more and more time on something, the more promising that it seems. So, we'll spend an initial quick one to three hours just checking out, "How many studies were done on this program? Are they high quality? Does it look plausibly cost-effective?" And if the answer to that is yes, then we spend sort of even more time drilling into our understanding, and then even more time... And to the point where we've probably spent, you know, 100 or more hours on it, we'll decide that this is a program that looks really good and we'd like to recommend a charity that's working in that space if there is a charity that then meets all our charity criteria. So that's our process.

**09:57 Speaker 2:** In regards to sort of metrics, how do you guys think about kind of indexing... Kind of have a broad range of potential problems or areas that we focus on to get donors to get a sense of, you know, impact per dollar on some sort of relative degree. Is that possible or is [10:23] \_\_\_\_\_.

**10:23 CH:** Yes, the question is, "How do we think about indexing or comparing programs in a variety of different spaces to one another. I'm going to hold your question, because the new research that I'm about to talk about is related to that question, so it is a really good question and it's one we think a lot about, so I'll share more on that in just a minute.

**10:41 Speaker 3:** Are there other organizations like GiveWell which measure specifically like cost-

effectiveness? Not... I know you said originally that there's some [10:49] \_\_\_\_ others, or... I don't know, I'm not familiar with that.

**10:57 CH:** Sure, yes. The question is, "Are there other organizations that measure cost-effectiveness?" There are other organizations in the global health space that calculate cost-effectiveness figures for various programs. There's various organizations that are focused on looking at the different costs of various health programs. I think something that we do that's unique is really focused on comparisons between programs that are really different, this question that we're going to get into in a moment, and coming up with the relative cost-effectiveness of those outcomes, and also really taking an all-in look at cost-effectiveness. So if you see a cost-effectiveness figure thrown out, let's say \$5 to save a life by buying an insecticide-treated net, that figure probably is just the cost of the net and doesn't take into account the fact that not everyone who receives an insecticide-treated net to sleep under to prevent malaria would have otherwise gotten malaria and died. It doesn't take into account the fact that the malaria rates in different areas can vary. It doesn't take into account the cost of distributing the net. So our cost-effectiveness estimates are really all-in and we try to think about as many factors as we can when we make them. But we're certainly also aware of other global health estimates and we'll look at them as we're looking at programs. I see a hand, yep. You, yep.

**12:16 Speaker 4:** How do you know balance a desire for detailed information about the charities that you audit with [12:22] \_\_\_\_?

**12:29 CH:** Yeah. So how do we balance, basically, wanting a lot of information from our charities with wanting them to be very cost-effective, and maybe our desire for lots of information could produce high costs to them, and how do we sort of weigh that. I think that we do ask for a significant amount of time of our top charities. When we're early stages in a review process with the charity we tend to try to ask them for maybe an hour long phone call, and then as many documents as they already have existing so that they're not creating new content for us to look at, so that we can then sift through them and if they seem promising kind of keep going. We offer charities incentive funding for participating in our process, so we insure that when you hit a middle of my milestone of a, we call it like an "interim charity review", you receive a grant of a \$100,000, and then when you become a top charity, where it's sort of the maximum time engagement with GiveWell, we make a grant of, last year it was, 2.5 million dollars, to each of our top charities, to incentivize them to engage with what definitely is an intense process.

**13:32 CH:** We also think that a lot of the information that we want is useful to the charities in maximizing their impact, because we're looking for information about how they're measuring that their programs are working and how are they monitoring to ensure that their beneficiaries are reached. And you could have an organization with a very low overhead cost that doesn't get any of that information, but then you would never know that their program is working. So we think that some of those pieces of information that do take staff time and cost money to collect are also just invaluable pieces of telling whether a charity is effective at its work.

**14:08 Speaker 5:** So is dollars per life saved like the main quantitative measures? Are there any other goals other than saving lives?

**14:18 CH:** Yes. Okay. "Is dollars per life saved the main thing that we look at or are there other goals?" I mentioned that we have five top charities that we recommend for their focus on increasing

incomes or consumption, as well as three that focus on saving lives. So those are the two main areas that our current top charities list works on. I think it's a good moment for me to shift into the next topic which talks about how we compare those areas, which is a very challenging question. So, without further ado, and there'll be more time for Q&A after that, I'm going to shift into talking about some new research that we've conducted, or that we've supported, on how the potential beneficiaries of our top charities think about the different good outcomes that our top charities achieve and how they would prioritize among them. So, GiveWell has a set amount of donations that we allocate each year from the community of donors who use our research. As I mentioned earlier, this amount isn't the full amount that we think our top charities can achieve, and so we do prioritize among our top charities to try to figure out where should we direct funding first, second, third, and so on.

**15:26 CH:** And because we think that increasing incomes and improving quality of life, as well as saving lives, are really important, we're trying to directly compare charities that work on those two different outcomes with one another. And this is really hard to do. We've looked into the existing literature to see how other people have studied this problem. There's some research that has been done in high-income countries on how much more people need to be paid to take on jobs with a higher risk of death. That's a common way that people sort of think about these income versus health tradeoffs. As far as we know there has never been a similar study conducted in the parts of the world where GiveWell's top charities work and that this is sort of a major gap in the existing research. And so, in the time before this new research that I'm about to talk about, we've tried to approach this problem by looking at what other major global policy makers do to make these tradeoffs. So, looking at groups like the World Health Organization, talking to philosophers, building our own models and surveying staff to try to get a sense of how we should think about these different good outcomes relative to one another.

**16:43 CH:** But we also thought that we should help create this type of information which we think would be really valuable. And so we supported a group called IDinsight, which is an organization that we've worked with over the years through a part of GiveWell's work called our "incubation grants program", where we're trying to conduct research and develop potential future top charities. We supported IDinsight to survey potential beneficiaries of GiveWell's top charities to see how they would value these different good outcomes relative to one another. So IDinsight conducted this survey between May and September of this year. They interviewed over 1800 individuals in low-income households in Kenya and Ghana and we very recently got the results of that report back. So at just sort of a very... Just to give an example of sort of what this research looks like, these are really hard questions to ask. IDinsight took three different approaches to asking these questions. One was the individual-focused approach, where as you see here on the screen... And this specific question is a slightly simplified version of how they asked it in the actual survey, but it's very similar to the content that was there.

**18:04 CH:** They asked basically how much would you be willing to pay for a vaccine or medicine that reduces your risk of contracting a fatal disease from 20 in 1000 to 15 in 1000. This is an approach that's aligned with the existing literature on how much people need to be paid to reduce their risk of death in various jobs, but obviously has a challenge where it's really hard to ask people about small probabilities. I don't know if any of you are looking at that question and trying to think about what you would answer. It's a really hard question to answer. The other challenge of asking the question in this way is that participants' responses are constrained by their income, or their ability to pay, since you're asking someone how much would you be willing to pay, it follows that

the amount is income constrained and not sort of the general amount that you might want to pay.

**18:56 CH:** So they also asked these questions in another way, which was putting the survey respondent in the position of a resource allocator and saying, "Imagine that you're in charge of the resources for this community. Here's two programs that have these different outcomes. Which one would you choose?" And this is an actual question, it's in quotes, this is as asked on the survey. So, just to give a sense of how these questions are asked. And this is basically asking how do you value the three GiveWell charities that focus on reducing the likelihood of children dying of malaria and of other infectious diseases and GiveWell's recommended charity GiveDirectly, which distributes cash transfers to very poor households.

**19:44 CH:** So there are challenges here. You don't have the challenge of being constrained by what you can pay, since you're asking them to be in the position of a resource allocator, but you might have challenges of social desirability bias where people's answers might be influenced by what they think is the right thing to say, that you might think that you would get more responses that lean toward health or life saving interventions for this reason. And the third approach besides these two was also that IDinsight, the surveyor, asked participants to walk through their reasoning for each of the decisions that they made, so they could get a qualitative sense of where the answers were coming from and that might help a bit with interpretation. So I have a note here to please interpret these findings with care. GiveWell is still going through the results of the study which we got back in the final form last week, and we also, for the reasons that I just mentioned, think that these are really hard questions for anyone to answer. We wouldn't want people to put too much stock in a single survey of 1800 people being asked these really hard questions.

**20:54 CH:** We broadly saw that the results sort of nudged in two different ways. One was putting a higher value on averting death over increasing consumption or ability to buy things, and also put a higher value relative to what GiveWell's past approach had been, on averting the deaths of children under five. We're not planning to make dramatic changes to our cost-effectiveness model this year. As I mentioned, we're still in the stages of reviewing the study and I think that these are results that you have to treat carefully. We hope that this inspires other people to contribute to literature here. We hope that people will take a variety of approaches to getting more information on these hard questions so that we can improve our understanding through having a robust data set. But we're really excited that this is a step in the direction of having more information to help us make these hard choices within our cost-effectiveness model. So I'll pause here for questions on this.

**21:54 Speaker 6:** How did you ensure that the respondents understood the question?

**22:00 CH:** Yeah. So, the question is, how did IDinsight ensure that the respondents understood the question. The respondents were each put through a sort of training module where they were asked about a series of small probability questions. So a type of question might be, "You're going down the road and you can choose to go down path A or path B and path A has a 10 in 1000 risk of death and path B has 20 in 1000 risk of death. Which do you choose?" So they would go through a training module first to kind of get a sense of whether respondents sort of understood the small probabilities. They also used visual aids in some context to show the difference, over here there's five lives and that's represented visually and over here there's six lives and that's represented visually. And I believe in most or all cases where there was a lot of inconsistency in the training module responses, that they didn't end up incorporating those into the final results. Yep.

**22:53 Speaker 7:** Either in this context or in the broader effectiveness models, does GiveWell think about the derivative economic effects of giving, either by saving for [23:08] \_\_\_\_\_ or deworming treatments, on not only on the recipient but on the broader foundation realm?

**23:13 CH:** Yes. The question is do we sort of model economic benefits beyond the individual recipients of our top charities? It depends on the charity. For GiveDirectly, which is the charity that we recommend that distributes cash grants to very poor households, we do look at the total size of the household and how many people are affected there. In other cases it might be too challenging to model what the community level effects are of increasing income. And similarly, it would just depend on the specifics of whether we felt we could reasonably model it. We don't really try to model factors that we think would be very important to the bottom line understanding of the program, and also factors that we can reasonably model. And if not, we sort of move them to either qualitative assessment or sort of unmodeled adjustment. So it would just vary depending on the charity, the type of data we have, and the type of program that they implement.

**24:06 Speaker 8:** Catherine, what do you mean by higher on this? So, is this something where it was disproportionate and there's obviously some noise or was it just statistically significantly higher?

**24:18 CH:** What do you mean by... Like placing more weight on this one?

**24:20 Speaker 8:** Yeah.

**24:21 CH:** Yes. So the results sort of pushed in this direction. I think it varied a lot, the magnitude, depending on the way that you ask the question. So I mentioned there were these two different ways to ask the question, either putting you in the position as an individual and saying, "How much would you be willing to pay?" or putting you in the position of a community resource allocator, and you actually got quite different results there. But the overall push of the report was in these directions somewhat clearly, but we're not sure how much we should update our cost-effectiveness based on them as we're still kind of going through and trying to think through how much weight to put on this versus some of the other things that we've looked into over the years. Yeah.

**25:03 Speaker 9:** I'm curious... This is kind of two questions. One of which is how the surveys are designed, like how much... Like what portion of it was GiveWell's and what portion of it was sort of IDinsight, but also how did you... How did IDinsight parse out respondents from participant countries? What was the approach for that?

**25:27 CH:** Yes. So the question is how were the questions designed and then also how did IDinsight find respondents, I think, is that... Yeah. So GiveWell provided research design input into this survey. IDinsight also did. So that was a collaborative effort between GiveWell and IDinsight where we spent a lot of time talking to them and thinking through what information we would need to get out of the survey in order for that to be useful to us and the types of questions that we try to answer in our work. They also provided input on what we're asking.

**26:03 CH:** And then to find beneficiaries, so we were looking for beneficiaries that lived in areas where GiveWell's top charities work and that also were in very low-income households, since that's the population that GiveWell's top charities tend to serve. And I believe that the way that they identified those respondents was by... It varied a little by country context, but they tried to find a



sort of four different geographies that represented somewhat diverse populations. And then within those geographies to randomly select villages to work in and then within those villages they took different approaches to identifying the poorest households. And that methodology is all in the survey responses that we're going to be publishing in a couple of weeks on our website, if you really want to dig into the details. But basically, they're trying to find populations that are really similar to the ones that our top charities serve. And then kind of randomly pick once they had nailed those down at a higher level.

**27:00 Speaker 10:** Are you asking open ended questions like, "What do you think would be the most effective way to help your family or your community", something like that?

**27:08 CH:** Do we ask any open-ended questions about the best way to help family or community? I'm not positive if they were asked of all recipients, they did do a number of more intensive qualitative interviews with some number of people to talk about their lives and ways that they thought would be most helpful. I don't know if there was an open-ended question included by default, but the qualitative reasoning portion... So, that I mentioned is kind of the third approach, where they asked each respondent to talk through their rationale for picking what they picked. So, they talked about some of the factors would be like, "Well I'm particularly interested in helping a young children, because they're the future of this community and I think if we're able to help them that's really great because they'll go on to help us all". There were also responses that focused on the benefits of receiving cash as kind of a kick-starter for the ability to have businesses or address health challenges. So they did get more open-ended responses for why people were making the choices between that. But I'm not sure how often the question of just probably speaking what would be the best thing to have happened was asked. Wait. Sorry do you have one?

**28:14 Speaker 11:** [28:14] \_\_\_\_ like, what you guys plan to do with these responses and what if they're not... Like what if they say, "Oh, we'll take the money, like, I don't care that much about kids' lives". [28:33] \_\_\_\_?

**28:40 CH:** Yes, the question is how do we plan to incorporate these into our research and what happens if they're really different from our own intuitions. The near-term plan is not to make major adjustments without spending a amount of time looking into the survey. I think we want to, check into how well we think the responses capture what our beneficiaries sort of truly believe. So looking in the survey methodology and all those questions. And then also thinking about how much this should be an input into our cost-effectiveness model where we put this in. So, I think it will depend on the product of much longer conversations that we've had at this point. And this year, we're not planning to make any major updates, but it's certainly something that we'll be writing about. You'll be able to see the full results of this survey on our website soon, and you'll able to see the way that they were incorporated into our cost-effectiveness model this year. Also very soon we publish, officially publish, our updated 2019 recommendations. Which I should say that this was a preview, we haven't officially published our 2019 recommendations. Yeah, some more information is coming. Maybe we have time for one more on this one and then we'll go on to Ben?

**29:49 Speaker 12:** Do you ever change your recommendations after you read [29:50] \_\_\_\_?

**29:56 CH:** Yes. The question is do we ever change our recommendations after we've made them. Absolutely. We are, as I mentioned earlier, in an ongoing review process with all of our top charities. Once they are on our list we continue to follow up to make sure we think that they're still

collecting good data to show that their programs are working, that they have additional funding needs.

**30:16 CH:** We've actually removed a top charity from our list in the past because nothing else changed other than we thought they didn't need any more funding at that time, so we removed them from our list because we didn't think they needed additional donations. So, yes, it's an ongoing review process, that's the sort of process we're figuring out if we should keep all of our top charities, we also have a research process going at all times looking for new potential top charities. So I talked about some of the program-level research we do, where we're looking at new academic research and trying to figure out if their studies that we should dig into and new program areas we should look at. That's happening in an ongoing basis, as is our research into new organizations that work on programs where we've already completed that process. So definitely always looking for new top charities. The process for becoming a GiveWell top charity is a pretty long road, usually hundreds, maybe more, hours of review. So it does take a little time from our initial contact with the charity to adding them to the list, but it's something that we're always looking to do and always excited to do.

**31:13 CH:** And I going to hop into the next section just so we have it have time for everything, there'll be more open Q&A time at the end, and then plenty of time to mingle and hang out after the fact. But just wanted to highlight a quick organizational update which ties into this question nicely. I wanted to share that our team has grown significantly this year. We hired 12 new staff members in 2019, and as you can see, they work on a variety of areas of GiveWell, but the overarching theme is that we're hoping to expand the research team to help us find more impactful opportunities, increase the impact of the money that we move.

**31:50 CH:** And then we've also expanded our outreach and marketing team to help us share GiveWell with a wider audience and hopefully direct more funding to the charities that we think are really excellent. So the team has grown, we now have 36 full-time employees at GiveWell and are excited to plan to continue our expansion into 2020. We're looking to hire senior fellows, applications for that are due very soon. But were looking... If you're a PhD Economist or know one we're looking to higher senior fellows with that background to help us look into new areas of research, particularly areas that we have historically found challenging to approach with GiveWell's traditional review process. So, looking into areas that are a little harder to measure, new spaces for us. We're also looking to hire on the outreach team a donor engagement specialists, and researchers at vary levels of seniority, and content editors. So a way that a lot of people hear about GiveWell is through people like you and the community of people who use our research. So, definitely would be glad for you to share these listings as we're hoping to continue our growth next year. As a quick personal update we moved to Oakland a few weeks ago. This is a picture of some of the team in our new Oakland office.

**33:06 CH:** And we're really excited to, you know, have grown so much that we no longer could fit in our old office and had to move across the Bay to a brand new office where we could all fit. So with this growth in mind I want to quickly turn it over to my colleague, Ben, who is our head of growth, and he's going to talk about some of the ways that you can help us during giving season this year.

**33:29 Ben Bateman:** Alright. Hi everyone, I'm Ben. It's great to meet you all, kind of.

**33:35 CH:** Here's the clicker.

**33:37 BB:** Oh, I have a clicker, great. So I wanted to talk a little bit about what we've done and then how you might be able to help us as we're growing. So, 2018 was our best year yet, we were able to direct more than \$140 million to our top charities. And this is really exciting and I want to thank everyone in this room who's a donor, or and gave and was part of that, and everyone else who's not and showed up anyways to hear us talk a lot about what we're up to. I think more than just that dollar figure is what we expect that's going to do.

**34:07 BB:** We think that would prevent over 30,000 deaths, provide parasitic worm treatments for over 70 million children, and provide cash transfers for over 10,000 families. Now I want to pause for a minute, because those are... Really easily can just blur into statistics, and I think it's really hard to think about what that actually is.

**34:24 BB:** You know, if you went home tonight and saved someone from a train as you're getting on the subway, I can't say, "Multiply that by 30,000 and that's how excited you should feel right now." [laughter] Because that's not how we work as people, but I think this is really exciting. I'm really inspired by this scale of impact that we're able to have. And I hope you all can take a minute, whether it's tonight, whether it's at another time, and think about your role in this and how exciting we should feel about it. I think it can feel very gated behind statistics and gated behind numbers that are hard to engage with compared to what you might see helping someone in the day-to-day, but the leverage that you can get when giving to really outstanding charities is really incredible.

**35:04 BB:** And so, to talk a little about why I'm so excited about this, I think some numbers are helpful. So I wanted to talk a little bit about the challenges that our top charities address. So we have five charities that work to help people who are living in really dire poverty. So the proportion of people who live in really extreme poverty, that's less than \$1.90 per day, has gotten a lot smaller as the world population has gotten a lot larger, but that's still 730 million people, 2x the US, who live in really extreme poverty.

**35:39 BB:** And everyone else is living in what would seem to be pretty poor, living here in the US. Most of the world lives on less than \$10 a day. And this isn't a problem we could solve all at once, but we do have charities that we recommend that we think do really outstanding work to help people who do live in extreme poverty. Our charities who do deworming and GiveDirectly who does cash transfers. I also want to talk a little bit about the number of people who are dying from malaria.

**36:05 BB:** Hundreds of thousands of people every year, and most of those children under five. And we have charities that do really amazing work here, Against Malaria and Malaria Consortium. We expect to be able to prevent someone from dying from malaria for under \$3000. So I'm sharing this context because I'm both really inspired by the amount of good that we're able... This group, people who are donating and people who are thinking about this issue, are able to do, but also this scope of the problem that we're facing on how big these challenges are. The charities we recommend do have giant funding gaps and we're thinking really hard about how we can [36:38] \_\_\_\_ influence to try to help address these problems and put more funds to do a really amazing amount of good for the charities we recommend.

**36:45 BB:** So there's some things you can do beyond giving that we think would be really helpful,

and that's sharing GiveWell with other people in your life. Now a lot of donors I know have come in because they are a friend or a family member who has gotten them excited about this. I am one of those donors. My friend, Chelsea who, caveat, was a GiveWell employee at the time [chuckle], talked to me a lot about giving, and I was in tech and had never thought about giving, not associated, but just separate. [laughter] But she recommended some books, and I read them and it led to some long conversations about what's meaningful in my life, and now I work with GiveWell. [laughter]

**37:20 BB:** That's not the arc of everybody that'll talk to you, but I think sharing more about effective giving can be really helpful, but can also be challenging. So a couple of things that we think are helpful, the first is just to share articles or books or media that can start a conversation, or just expose someone to these ideas. So we've actually put together a list of things that we think are good introduction points on our website on [givewell.org/howtohelp](http://givewell.org/howtohelp). So you can look these up and you can share those with friends and that might lead to them reading about what we're up to or just start a conversation if you want to follow up there.

**37:52 BB:** And you can also invite us to come and speak at your workplace. So we've done some workplace talks and we know that learning about effective giving from co-workers can be a really great way to get exposed to it, and it's really low effort for you, where you can bring us into your work and you don't have to go and have all of those conversations yourself. So if you're interested in that you can go on our website, and you can sign up. We have myself in New York, and we have a team in San Francisco where it's very easy for us to come and speak at someone's work, and then we can go further afield if we know there's a few companies in that same place, and that's worked well for us in the past. Is that it? Thank you so much for coming. And then Kat and I are both around to answer questions if you have any.

[applause]

**38:40 CH:** So we have a few more minutes here of open Q&A where we're all together, and then we'll take a break and mingle. There's still some drinks left and food, but let's hang out here for a few more minutes. Any questions on anything left from the conversation? Yeah.

**38:53 Speaker 13:** I have a question about the overall strategy that GiveWell uses to first identify the new charities. So, hypothetically let's say there was a charity that had... Or a problem that had a higher [39:04] \_\_\_ burden of disease, compared to something that's already working. So, let's say even though there's a charity that's struggling to find funders to help them, even that they have a higher burden of disease, let's say, you can take an example of [39:17] \_\_\_, there's a dozen or more other [39:20] \_\_\_. And there might be other diseases that have a higher burden, but maybe there are charities that are working on [39:29] \_\_\_ working as well. So does GiveWell think about the burden of disease when starting with [39:34] \_\_\_ starting with the question, or does it consider cost-effectiveness and what programs are working where?

**39:39 CH:** Yeah. So the question is, I think, how we would trade off an organization that works on a program that's maybe not the most pressing in terms of its burden, but they're excellent and they work on this problem and exist; versus trying to find a charity that works on a very pressing problem; or maybe finding a charity there that's not as effective, but the thing that they're working on is more important in terms of how much impact it has on the world.

**40:05 CH:** Yeah, so in terms of the way that this intersects with our research process. So, I mentioned the first step in our process is trying to identify the program areas where we think we can have the biggest impact and how significant the burden of the diseases in the world might be one of those impacts. So, malaria is a good example where there are many people affected by malaria. We also know that it's a very cost-effective disease to work on because there are cheap ways to prevent malaria. So that kind of ticks both of those boxes maybe, and is easy.

**40:36 CH:** Other times we will find programs or maybe there is a very cost-effective treatment, but the burden isn't high. There we might still be interested in providing funding if we think that that's something that is possible. So the thing that I'm thinking of here is maternal syphilis. So there are pregnant women that have syphilis that can pass it on to their children, and that could have serious health consequences for babies or fetuses, and the total global burden is actually not that large. I think it's about 1.4 million pregnant women in the world have syphilis, but the treatment is very cheap and it's also quite neglected.

**41:20 CH:** So there aren't charities or funders that really focus on this problem perhaps because it's small, it's not very well-known. We've been interested in whether we might be able to fund the scale of a program to provide those treatments. So even though it's not a huge, huge area, we still think that that is really important, because it's very cost-effective and seems like a well evidence-backed intervention. So there are kind of many different inputs and it depends a little bit on, I suppose, if there was a program that we thought was very excellent, but didn't need any more funding... That will sometimes come up.

**41:55 CH:** We've struggled with, in the past, recommending charities that were on vaccines, which we think are very... There's a lot of evidence that there are very cost-effective ways to improve lives and increase people's health, but that there are also many funders that work in that space, so we haven't recommended vaccine charities, historically, for that reason. So those are all factors that we take into account when we're thinking about whether we should work in that space.

**42:23 Speaker 16:** So, now I mean, of course, that your work with charities is very different than the sort of for-profit investing space. But there's movements in like [42:31] \_\_\_ investing, impact investing and the sort of lot of mind investment, and a push towards greater methods of... Greater levels of assessment and increasing focus on methodology for impact assessment in that area. I'm wondering if you think about your work in relation to that at all? Or if there might be any overlap between methodology and what you're doing, and what other communities might be trying to do?

**42:57 CH:** Yeah. So the question is, do we see intersection between our work and impact investing, or for-profits, that are looking into sort of double bottom lines. I think we are open for our methodology to be shared with anyone, and we think it's a really good way for assessing impact. We would also be open to recommending something if it made sense in the for-profit space, but our guess is that a lot of the time the focus on the bottom line would lead you to be less effective on either bottom line. So that we think our charity is really focusing on the things that they focus on, leads them to be extremely cost-effective in what they do. We have done some work on results-based financing as an interesting way to intersect with how major aid agencies are thinking about their funding. So that's kind of in the related space, something that we're thinking about, but generally we haven't sort of set a major strategy focus on impact investing as an area.

**44:01 CH:** Way in the back.

**44:02 Speaker 14:** Because you tend to recommend the same top charities every year, do you worry about diminishing returns on conducting evaluations on the same charities every year?

**44:16 CH:** Yeah. Because we recommend so many of the same groups every year, do we worry about diminishing returns to conducting our evaluations of those groups? Yeah, I think it's a really good question. Because a lot of our review process is not just centered on whether a charity qualifies as a GiveWell top charity, but also how we would prioritize it in our list of top charities, we actually see a lot of movement year to year on the relative cost-effectiveness of our top charities and their ability to use additional funding well.

**44:44 CH:** So a charity in one year might have four new opportunities to expand in new countries, and we think that there's a huge upside to them trying these new things and these countries have, let's say, huge burdens of the disease that the charities are focused on, and these look really cost-effective. And then the next year, they've kind of already moved into those spaces, they don't have as any opportunities, they're kind of going to the next countries which maybe have a lower disease burden, so they're less cost-effective, so then they get sort of moved down our rankings.

**45:07 CH:** So what we really hope is that people won't just focus on the eight groups on our list, but really the relative ordering of those groups, which is a major piece of our research work every year. We really think hard about which of our top charities needs the funding in the most pressing way. But on an other scale, we don't prioritize asking each top charity all of the questions they've been asked before, every year. As I mentioned, we do try to focus on the updates that we think will move our understanding the most, and so we'll be really looking for prioritizing those pieces of information each year and kind of asking ourselves, "What would change our recommendation of AMF this year if we learn that?" "What would change our recommendation of give directly?" And trying to narrow in on those pieces of info.

**45:56 Speaker 15:** I imagine in a lot of the research you guys do, you feel research constraint for your goal of being able to figure out relative cost-effectiveness both on even top charities like the deworming question, but also all kinds of other spaces where I feel like [46:14] \_\_\_\_ sometimes says we just don't feel like there's been enough research to support this compared to malaria or compared to [46:24] \_\_\_\_ . Does that ever make or has it made GiveWell think about the possibility that the most effective use of funds is trying to get really high quality [46:34] \_\_\_\_ on some of these less researched spaces, or spaces that need more research like deworming for example.

**46:42 CH:** Yeah, so the question is, doesn't it seem like maybe one of the most impactful things we could be doing is funding research to help us identify really effective things, and we actually have the same thought. So we have funded some research into areas that seem really promising as part of our incubation grants program, which I just very briefly mentioned, is a part of our work where we're trying to build research to support and help organization scale that we think might become top charities in the future. So we have provided funding for a few RCTs that way, also research like the beneficiary preferences survey that I spoke about tonight, which was a direct response to us, seeing a gap in the literature, that we were hoping to fill.

**47:23 CH:** And so currently, I think that that type of work would probably live in our incubation grants program, which is funded by one of our long-time donors an organization called Good Ventures. So it isn't something that we're sort of actively directing donors to, but we think might be

a really impactful thing to support that we hope will lead to the creation of new top charities.

**47:46 Speaker 16:** I'd be interested in each of your answers for this. If you were going to guess, where really left field charity or cause area was going to your recommended list in say the next five years, where will it come?

**48:01 CH:** I can go. I guess I'll give my best take. I don't know if this counts as left field, because we've been talking about it a lot, but I think the thing that would be a big shift for us, or a new area for us, is a charity that focuses on policy interventions. So one of the areas in which we're expanding our research into is looking into supporting policies in low and middle income countries that might potentially affect many people, in a positive way. And so, you know, if you can imagine like for a relatively limited amount of funding supporting an organization that causes a policy to be passed, or causes a policy to be enacted very effectively, that that can be a very cost-effective thing that you could do that could have impacts for a long time. So not fully left field because we have been looking into this, but I think would probably feel it maybe a bigger departure from here than some of the other things that we've talked about looking into.

**48:56 Speaker 17:** Ben you have to say something different on this.

**49:02 BB:** An intervention that we've identified that there hasn't been a great charity for, seems like the best candidate, I don't think at that left field, but that would be my best guess.

**49:13 CH:** Okay, so someone who hasn't asked a question yet.

**49:15 Speaker 18:** If you had to make a guess about what is the probability of zero impact for your top aid charities, like what were [49:21] \_\_\_\_ your confidence interval around nonzero impact is?

**49:28 CH:** Yeah, so what's the best guess of the probability of zero impact of one of our top charities.

**49:31 Speaker 18:** Yeah, how confident are you that they are doing something?

**49:33 CH:** Yeah, we don't... We do not publish confidence intervals because the way that we do cost-effectiveness, it doesn't end up making sense within our model. I think the intervention that we've written the most about this potential for is deworming. We actually have a blog post called deworming may have huge impact or may have zero impact on our blog. We don't think that that means that deworming treatments don't treat worms, we think that they do treat worms and there's good evidence that they're effective in that way, but the longer term picture is that we think that children who receive deworming treatments may go on to earn higher incomes when they enter the workforce as adults.

**50:14 CH:** To make that recommendation we rely a lot on one study that was done in Kenya in the late '90s, and a series of follow-ups to that study, we have spent a very large amount of research staff capacity investigating this study, we've commissioned additional reviews of that study, and spend a lot of time thinking about how likely this one study would be to actually hold up today. And we feel good recommending deworming, but I think that's the one where we have the sort of smallest evidence-based in the sense of number of studies to rely on compared to malaria where we have significantly more randomized control trials of the malaria programs that we recommend.

**50:55 CH:** Deworming is my best guess that fits that question, but one that we have also spent a lot of time on and feel good recommending. And it's unfortunately, 8 o'clock on the nose, so we're going to end the formal presentation, but we are all still here ready to talk about GiveWell, all the GiveWell staff. We also want to thank Will for providing beverages for tonight's event, so we hope that everyone should give him a round of applause. We hope you'll hang out, have a glass of wine, continue the conversation, and again, really appreciate you supporting our charities, coming here and supporting our work, being interested in hearing about our new research. It's really great to see you all here and to meet everyone here, so thank you.