

GiveWell NYC Research Event, December 14, 2015 –  
Top Charities

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

Elie: All right. Well, thanks everyone for coming. I'm Elie Hassenfeld. I'm one of GiveWell's co-founders. This is Natalie Crispin.

Natalie: Hi. I'm a senior research analyst at GiveWell.

Elie: We're really happy that you joined us tonight. I just want to quickly go through some of the logistics and the basic plan for the evening. Then I'll turn it over to Natalie to talk a little bit. Just first, like we do with pretty much everything that GiveWell does, we're going to try to share as much of this evening as we can with our audience. That means that we're going to record the audio of the session and publish a transcript on our website. If you say anything that you would prefer not be included in the audio or transcript, just email me, [Elie@Givewell.org](mailto:Elie@Givewell.org) or [ingo@Givewell.org](mailto:ingo@Givewell.org) or just come find me at the end of the night, and we can make sure to cut that from the event. You should feel free to speak totally freely this evening. We can deal with the transparency part afterwards.

The basic plan is that the first thing I'm going to talk about is GiveWell's top charities. It's December, and we just refreshed our recommendations. Natalie's going to go through the strengths and weaknesses of our top charities and some of the major changes over the past year. Then we're going to take a little break for questions and we're going to break periodically throughout the evening to give you a chance to ask whatever's on your mind. Generally speaking I would encourage you to just ask whatever. Don't try to make sure you're only asking something that we have not written about. One of the main reasons that we do events like this is that we want to enable people to ask questions that they might not have had a chance to find the answers to on our website. We really hope you'll take this chance to ask us whatever's on your mind regarding our work.

Regardless, Natalie's going to start talking about the strengths and weaknesses of our top charities. We'll break for questions. I'll then talk a little bit about the future of GiveWell's research and some of the growth that we've had over the last few years. Then we'll take a brief intermission so people can get some more food. Then we'll come back and talk about the open philanthropy project, which is another part of our research. First I'll give a high-level overview of that project. Then Chloe Coburn, who is our program officer for criminal justice reform, she's over there. She'll go into the more detail on her work and her efforts in that area. Without further ado, Natalie, why don't you take it off?

Natalie: Thanks everyone for coming. Really great to have this many people to talk about our work with. I've been at GiveWell for 6 years, and I work on our top charity recommendations. Both following organizations that we have recommended in the past

and finding new organizations to potentially add to that list. I'll give a brief summary of what has happened this year, top charity strengths and weaknesses and try to just give some basics so that then you can ask questions and let me know what you are more interested in knowing more about.

First, the process for finding our top charities. Very brief. It starts with an assessment of the academic literature. That focuses on the program level. Vaccinations or deworming or bed nets to protect against malaria. We do that because in our experience, charities don't have the type of information that would lead us to be able to establish the cause and effect of a program. The type of information that we often rely on are randomized control trials. It's not the only evidence we look at, but it does form a big part of those assessments. We then use that information to look for organizations that are running programs that we've identified as particularly cost effective and are implementing those programs well.

In 2015, we largely looked for organizations that were working on priority programs. That full list is on our website, but it generally includes global health and nutrition organizations as well as unconditional and conditional cash transfers. Other examples are micro-nutrient programs, vaccinations, and then the programs that our top charities are running. Bed nets, deworming. Overall, you may have seen that our top charities are the same as they were in 2014. Their strengths and weaknesses really briefly. There's AMF and Deworm the World Initiative. We think both of those are very high quality organizations. AMF, sorry. Against Malaria Foundation. They are running bed net distribution programs to protect against malaria and save children's lives.

The Deworm the World Initiative provides assistance to deworming programs to deliver pills to kill intestinal parasites. Those two organizations are running very cost effective programs. We think they're high quality organizations. The Against Malaria Foundation, in the past, we had a major concern with them, that they were not moving money out the door and delivering bed nets as quickly as we might have expected. They had a limited track record of working with partners to carry out very large scale distributions compared to what they had done in the past. In 2015, we saw an improvement in that track record. They went from working with one partner in Malawi to working with two partners. One also in the Democratic Republic of the Congo. It's still a fairly limited track record but a big step for them. They were able to collect the monitoring data they aimed to.

There were three major updates on Deworm the World Initiative. One was that they worked with the Indian Government to establish a national deworming day. Before that, individual states might carry out deworming, but there wasn't a national push for it. We are not confident that they played a key role in that, but think there's a possibility that they did. If they did help expand the program to millions of more children. They also improved the quality of the monitoring that they collected and shared with us so that we have higher confidence in that deworming treatments are getting to the children that they aim to get to and the quality of the program is generally high.

They also have significantly more room for more funding this year. They had several opportunities to expand to new countries. We think that they will continue to have

those opportunities. The third organization is Give Directly. They fund unconditional track cash transfers to households in Uganda and Kenya. They are a very high quality organization that does a somewhat less cost effective program. We think AMF and deworm the world are higher cost effective programs that we think Give Directly is a real high quality organization, and there's a very strong link between what they do and impact on the ground.

Overall, we think that their operations continue to be high quality and were expanding in both countries that they were working in. They also put additional effort into networking with potential partners. One such partnership led to a program that's going to start next year with a bilateral aid donor and funding from Google in order to text cash transfers against a more traditional aid program and see which is more effective. The last top charity is the [inaudible 00:08:12] Control Initiative. They, like Deworm the World, fund deworming programs. We think it's a very cost effective program. We had in past years and have continued to have some communication challenges with them and generally have not been able to have a concrete sense of how they have spent funds and what funding they hold. That continues to be the case. The other major update with them is that they provide additional monitoring that we thought somewhat increased the case for their impact.

We also continue to designate four standout organizations, and those are the same this year as they were in previous years. The Iodine Global Network and Living Goods. There are no real major updates there. We did continue to speak with them. There were some significant updates on Development Media International and Gain. Development Media International produces radio and television programming to encourage people in the developing world to change their health behaviors with the aim of saving children's lives. Last year we had reviewed the mid-line results of their trial to test how effective their program is. Those results showed increases in the behaviors they were targeting in some of the behaviors they were targeting.

They then came out with end-line results this year that they shared with us. Those results were not as positive. They did not show changes in the behaviors, and they did not find a decrease in childhood mortality, which was the aim of the program. There's also the Global Alliance for Improved Nutrition. The program we looked at with them was Universal Salt Iodization Program. They work with governments to try to increase the access to iodized salt. We spent significantly more time in 2015 trying to understand their program. We spoke with their staff in India as well as sent two staff members to Ethiopia to learn more about the program there. We ultimately found that we were not able to establish a clear case for impact for their program, and we attribute this to difficulty in connecting their activities with the impact that they aim to achieve as well as communication travel difficulties that we've had with them.

Those are the four standouts. In terms of what's changed overall with our top charity recommendations, we believe there's a stronger case for our top charities this year than there was last year. We think that the organizations have a longer track record. Just an extra year of us following them and them operating. Those track records have generally shown success in operating their programs. We've also seen in some cases where there's been more in improved monitoring. Particularly with the Deworm the World

Initiative. Then there's also significantly more room for more funding across the board. Each of our top charities have more opportunities to spend funds productively, we believe. Commensurate with both those considerations, we recommended more funding from Good Ventures, the foundation that we worked closely with, which was started by Dustin Moskovitz, the co-founder of Facebook and Asana and his wife Cari Tuna.

Together with recommendations from earlier this year, they're giving 44 million dollars to our top charities. We also think that the amount from individuals is likely to grow this year. We're projecting about double over last year. All around, significantly more money moved to our top charities. Because of both of these things, both more ability for the charities to absorb money productively and more money going to them, we've spent significantly more time in thought into this concept of room for more funding. How to rank relative opportunities against each other. Thinking more about what constitutes a particularly promising funding gap that we want filled in, ranking those funding gaps against each other. We've written a bunch about this on our blog. I won't go into details right now. Please feel free to ask questions. If there's particular aspects that you'd like to know more about, I'm happy to talk about that.

Let me pause there. Thrown a bunch of brief updates at you and interested to know what you guys want to hear more about. Please do ask. Yeah.

Speaker 3: Speak a little more to how you really know for sure that the, because I think you were probably very thorough on that, a little more about how you go about confirming that they did what they're recording really happened and how you know that they're more effective than similar organizations [inaudible 00:13:17].

Natalie: Sure. For each of our top charities, we've looked in depth with the monitoring information that they've shared. They each do it in a different way, and we don't just take them at their word for what they say they have available. We look at the data ourselves and try to find discrepancies and ways to try to understand the methodology really thoroughly of what they've done. For example, the Against Malaria Foundation does post-distribution surveys, 6 months, 12 months. Every 6 months for two years or so afterwards, they do these surveys where they go house to house and check whether bed nets are in place and being used on a regular basis. We've looked at the results of those as well as the reports on how they carry those out. Each of our top charities have processes like that to not just assume that the program has worked but to really check that it's happening as planned.

Elie: Just to add a little bit to what Natalie said, it's definitely possible that a charity is sending us fraudulent materials. We do not have a process right now for auditing those materials to determine if they're trying to defraud us. I would say a lot of our goal at this point has been forced charities to move the standard. Tell a story and get a donor excited to the point where they would have to be actively producing fraudulent materials in order to get a GiveWell recommendation. For them to convince us that we should support them. It's interesting because this is something that was a really good answer to this question 4 or 5 years ago when the amount of money we were directing

was less than a million dollars. I think it becomes less of a good answer as the money we direct grows in the incentives to potentially do this increase as well.

One of the things that we've been talking about as we plan for future years for 2016 is whether we should be trying to do more, given the amounts of money that we are influencing, to try and look for even bigger problems, look for potential fraud or potential charities. Given how close we know these organizations, I find it very hard to imagine that this is happening, but nonetheless, I think it's a reasonable thing for us to do. The other part of your question. I should add one thing. One small thing we've done is we've hired a journalist to go and visit a couple of our top charities. This journalist, Jacob Kushner, went to Kenya to visit both Give Directly and Deworm the World and see whether anything looked really problematic. He ended up reporting things that the charities both took to be pretty negative but we took to be pretty positive.

Because what he reported was, and you can read about this on our blog if you Google Jacob Kushner, Givewell, you'll find the blog posts. He reported some degree of displeasure in the villages that Give Directly had given money to among people that hadn't received the money or parents of children who were dewormed in deworming programs in schools that they didn't feel fully aware of exactly what the program was. I think these were legitimate issues that we have written about on our website, but they are far from the concern. The program wasn't having the effect that we hoped at all.

That's about the fraud. The second question is how do we know that the Against Malaria Foundation, which distributes bed nets, is better than Nothing But Nets, a very well known organization that funds bed net distributions. There, the true answer is we don't know. We don't have any information that tells us that Nothing But Nets is doing a bad job. What we know is that the Against Malaria Foundation is able to share significantly more information showing that its programs are successful than we've been able to get from Nothing But Nets. I think that groups that are willing to share, maybe people have different intuitions, but I put some weight on that desire to be transparent as an indication of a better run program. In the case of nothing. They've declined to participate in our process for a couple of years. It definitely does not mean that I think their programs are failing. If I had to bet which one seems like a better use of funds, I would give to the group that's sharing information and demonstrating its impact.

Speaker 3: I've worked [inaudible 00:17:58] non profit called the End Poverty Campaign. What we proposed is a comprehensive aid program where we'd like to see the U.S. government expand the Peace Corps and send representatives into individual villages and regions and organize the local people to monitor their own aid. Is that an idea that your organization would be especially when you have access to potentially major donors? I really appreciate your honesty in what you just said. It sounds like you made some good steps forward. Perhaps you might consider this to be a further step forward is having representatives from your organization or working with the Peace Corps go into the local communities and organize the local people that you have a direct link and you bypass all the bureaucracy and the potentials for corruption and siphoning off the resources. Is that clear, what I'm describing?

Elie: Yeah. One of the key criteria that we rely on is this criterion of evidence of effectiveness. Is there evidence that this program has significant impact relative to other opportunities? With something like malaria nets, there are more than 20 randomized control trials that show that if you give out nets, people use them and they save lives. We would then say does that type of evidence exist? Almost certainly it does not. We know, just to be clear, that when you have a criterion of evidence of effectiveness, there are things that are hard to measure, and you potentially leave really great opportunities on the table. GiveWell is trying to be a product for donors that meets the criteria that an outside observer can understand what we're doing, can vet it themselves and be confident that the donations are effective. Even if that means sometimes leaving really great opportunities on the table.

The other part of our research, the Open Philanthropy Project that we'll talk about more later, is the area where we're looking for more of those higher risk and potentially higher reward opportunities. Yeah?

Speaker 4: You mentioned that you often start with academic sources to determine cause and effect? I was wondering if that's like skimming journals or looking at university websites or personal contact?

Natalie: We use different methods, but I'd say the most common thing is to use published studies and, if possible, [inaudible 00:20:39] analyses that combine published studies in a highly rigorous way. We also reach out to experts when we have questions that can't be answered by that type of research. I think those are the main tools that we use.

Speaker 5: I find it interesting that you guys pretty much stuck with the same charities this year for both recommended [inaudible 00:21:08]. I was wondering if you saw any change in behavior of other charities and your relationship towards them? that didn't get them into those groups, but do you see any more [inaudible 00:21:19]? Do you see a trend of charities being more open to GiveWell and trying to provide more [inaudible 00:21:25] or is it really mean you're just coming [inaudible 00:21:27]?

Natalie: We reached out to several additional organizations this past year and got to an intermediate stage with them but didn't progress fully to be able to make a top charity recommendation. I'd say that some of those groups are potentially ones that would not have been interested earlier in our history. It's hard to know. While we have an open door policy of any organization can apply, we try to make our criteria as transparent as possible. I think we don't get a large number of applications to that process because it's a fairly stringent set of criteria. It's hard to know what the trends are in terms of transparency.

Elie: I would say as a general rule as the amount of money we are directing has grown, groups have been more willing to engage in our process. In early years, we would call up charities and they wouldn't want to talk to us, because the incentive that they had was very low, and now the incentive is much greater. I think one fundamental shift is that if there is a charity that we would really like to evaluate, by and large, they will spend time talking to us. Nothing But Nets is certainly an outlier in the respect that they have

chosen not to. That's been a pretty big shift. Like Natalie said, it's harder not to say. I don't think we have had any noticeable effect on the sector as a whole.

Natalie: Yeah.

Speaker 6: I'm [inaudible 00:23:05] question. I'm wondering have you guys found the effectiveness of a charity versus their stated initiative? It's like if you had a charity that had [inaudible 00:23:14] not as direct [inaudible 00:23:20]?

Natalie: Do you mean the they are? If an organization has strong evidence of effectiveness, what their stated goals are doesn't matter. I think in practice the organizations that we have been able to evaluate and establish a strong [case 00:23:45] for the evidence of effectiveness are fairly direct programs.

Speaker 6: What about Saving Lives?

Natalie: Saving Lives, Improving Lives, but more so just the number of steps between their activities and the ultimate outcome is somewhat shorter than say, an advocacy organization. It's just more difficult to evaluate that type of organization. That said, we've moved over time to be able to take into account more complicated causal chains. At first, we were largely focused on organizations that delivered products that saves lives or improved lives. The Deworm the World Initiative works with governments to deliver deworming programs. They do not directly carry out the programs or, in many cases, by the supplies. There's definitely a longer case in that chain and less direct impact. We think they are an outstanding organization and one of our top charities. The iodine charities that are on our standout list also are either advocacy or technical assistance organization. I think it's more difficult and has there's often less evidence for that. Our processes have improved over time and are more able to take that into account than they were in the past. It's definitely something we still struggle with and are trying to learn how to do.

Elie: Another way of answering your question maybe is, on some level, the goal is independent of our research process in that we often say we're cause agnostic, meaning, whatever activity you're trying to implement, we are willing to consider it. At the same time, our goal is to identify the small set of groups that we're going to recommend to donors. That definitely takes into account our judgement about the overall impact of achieving that goal. If a charity can demonstrate high effectiveness at doing something that seems less impactful to us, that will end up harming that organization in our research process. Just to put this very concretely. We don't only care about saving lives. It's saving lives, also improving lives. Right now, Give Directly, this group that tries to give 90 percent of the money they receive directly to very poor people, we think they are the most effective organization that we've ever seen at doing what they set out to, which is getting money into the hands of very poor people.

As a staff collectively, we would guess that bed nets are ten times as cost effective per dollar donated as getting cash into the hands of very poor people. Therefore, we end up ranking the bed net organization higher than Give Directly, because of that overall



impact. Just to be clear, we know that these cost effective estimates are very thorny and rely on a lot of difficult assumptions. Including this impossible to judge, philosophical judgement about how to wage saving a life versus improving a life. There's no question that is something one has to take into account. If someone feels very differently than us about that philosophical judgement, they should come to a different conclusion about where to give. That's something that we totally agree with.

Regardless, the goal does end up playing a pretty big role in the ultimate ranking of charities.

Speaker 5: Sort of a followup. In terms of your using words effectiveness and impact, high quality. Those are big with Give Directly. Translate those into something a little more tangible? Like effectiveness is something about reaching your goals and impact is the qualities of your light years. Can you make that a little more concrete?

Natalie: Sure. I probably have not used them fully consistently, but I think in concept, effectiveness is, yeah, how well you achieve your goal. The impact, I think I usually use it to mean what is the ... I guess often that takes into account cost effectiveness more. It's like how much good are you doing? How many lives-

Speaker 5: Per dollar.

Natalie: Saved or improved. Sorry?

Speaker 5: Per dollar.

Natalie: Per dollar, yes. Then what was the third one?

Speaker 5: You said particularly with Give Directly, [inaudible 00:28:34] high quality [inaudible 00:28:35] on that.

Natalie: Right. Yeah, I guess that takes into account both the effectiveness of the organization and achieving their goal but also transparency, communication, desire to learn from their mistakes. I guess those are a lot of the main factors.

Elie: Maybe to add to that a little bit. First of all, in a blog post that we wrote when we announced our top charities, there's a summary table. It goes through each organization on these different characteristics to just say where we rank them. Just so you know that exists, if you're interested. Evidence of effectiveness, I would say is probability that the program actually has an effect. Largely, we're relying on independent randomized controlled trials of that program. Randomized controlled trials are the medical style of assessing whether a particular intervention has an effect, where you randomly put people in a control group, in a treatment group and you see what happens. The only difference between the two, presumably, is the program itself. Often, we're relying on that for evidence of effectiveness along with the organization's own internal monitoring that we try to vet ourselves. That stands a very wide spectrum of quality from

organization to organization. There's really no rule of thumb about what that looks like on any individual basis.

Then there's impact per dollar. That's our best guess. What you get for each dollar that you give. In some ways, that is the only thing that matters, because all a donor should care about I guess is how much impact they get per dollar, but because this explicit modeling of that factor relies on so much guess work, we don't want to let that one actor drive the entire analysis that we do. You certainly could try to bake every single thing into one model, but if anyone's played around with these Excel models, which is what we use. You certainly could play around with ours. It's on our website. You just get in these situations. Is that number a 2 or is that number a 5? You don't really know. All of a sudden you've changed the cost effectiveness by a factor of 2 and a half. That doesn't seem like a very [robust 00:30:51] way to make a decision.

Then finally with organizational quality, one way we think about it is just, in a sense, it's this very subjective, like how confident are you in the organization and its people? I think that is a very concerning type of judgement for you all to be outsourcing to us and a very difficult judgement for us to be making. What we tend to rely on are groups that can answer our questions well. Where it seems that we're asking questions about program effectiveness. If they have thought about it before and have a good answer, we feel pretty good. Another way we think about this is would we likely know if some problem arose? In Give Directly's case, they found that a year and a half ago. Some of their staff in Uganda perpetrated a fraud, where they stole 20,000 dollars from the organization. That's a relatively small amount. They're moving about 25 million dollars a year right now. A small portion, but nonetheless, a meaningful one. They found it. They told us about it. They wrote about it publicly. That's the type of outcome that we expect from Give Directly, and we wouldn't necessarily expect from an organization which we've had more trouble communicating.

Natalie: Yeah.

Speaker 7: The menu or list of your favorite charities I think was very impressive in the methodology. [inaudible 00:32:15] deviations [inaudible 00:32:18]. I'm wondering why are there not opportunities in other sectors? I'm thinking here conservation of nature, historic preservation. In which there are not opportunities equally certifiably as rewarding as the ones you've picked in the deviation of suffering in Africa and India. I can think of a few quick hits in Seagrass Plantations or Reef Restorations or Saving the Architect [inaudible 00:32:54] in Havana, which could give you, I think comparably satisfying confidence in telling a donor that, yeah, this is really going to work and this going to be effective. [inaudible 00:33:13] these organizations that are doing [inaudible 00:33:15].

Natalie: We're coming to this work with a set of values that we are trying to accomplish. Those are not values that are shared by everyone. Definitely people could disagree with them reasonably. Our values include global humanitarianism. We believe that lives around the world have equal value. I guess that is the main one to your point. We wouldn't value historical preservation as highly because we would guess that it does not have the same effect on human suffering and achievement.

- Speaker 7: Do you think your methodology might be applicable, though?
- Natalie: Conceptually, yes, but I think in practicality I would guess that there are significantly fewer pieces of evidence to support the case for historical preservation. I have really not looked into that at all.
- Elie: One way you could [read 00:34:32] sometimes think about our work is that one of the questions we got all the time when GiveWell was starting was, "How does this scale?" People said, "There's so many charities out there. How can you ever look at all of them?" Clearly, we can't. We don't intend to. But because we make all our research public and we're trying to show what we do, we hope that others will take up the type of research we do. With historical preservation, it clearly wouldn't be exactly the same type of analysis, but the attitude of critical assessment of opportunities to find out [inaudible 00:35:09] giving opportunities could exist across many causes.
- We've actually seen one case of this happen already. There's a group called Animal Charity Evaluators. They're part of this effective altruism movement of which GiveWell is sort of a part. They are taking the GiveWell methodologies, so to speak, and applying it to look for groups that are trying to alleviate the suffering of animals and industrialized agriculture. I think for people with a certain set of values, that is potentially an incredibly cost effective cause to be working in. We hope that other types of initiatives like this will spring up, where people take something like the GiveWell approach and apply it to their area of interest. Hopefully that could help more people.
- Speaker 7: My question's [inaudible 00:35:57] thank you very much.
- Speaker 8: The list of top charities and standout charities seems to be a little more static the past couple years than it has been in the past. Do you consider this to be a good thing, because you think that you've stabilized, found the best charities or a bad thing because you're not uncovering further charities or [inaudible 00:36:21] or how do you view that?
- Natalie: We would definitely like to be finding great new charities that were better or comparable with our current top charities, and we continue to do that work. We, I think, didn't have enough capacity this year to do that as fully as we had hoped. I think that's part of the static nature. On the other hand, we have been at this for many years now, and we think there's a decent chance we've uncovered I guess the best ones with the methodology that we're using. We're constantly changing the methodology and improving it. I think there's definitely opportunities for us to uncover new ones and we hope to do that in the next year.
- Elie: It's at least a little bit of both. I think it's possible that bed nets and deworming in particular are such unbelievable giving opportunities that they're very difficult to beat by other interventions. Bed nets, it costs 5 dollars to get a bed net to cover 2 people for 2 years. Estimating it all out, it's 3,000 dollars per life saved. That's our best guess. 3,000 to save a life seems pretty crazy as something that you could actually take money and do. Deworming is a 50 cent pill that if you give it to a child through their middle school years, 10 years later, the study shows significant increases in income.

There's a ton of problems with this evidence. That would be a whole 'nother conversation if people are interested. Nonetheless, it's a 50 cent pill that treats and clears a parasitic infection. There's some evidence that this might have astounding impacts on long-term development. I think these are 2 opportunities that I am both hopeful and expecting that 10 years from now, these types of giving opportunities won't exist anymore, because I would just guess that donors will have closed these gaps. I think in total there's something like on the order of half a billion dollars of room in these interventions. 500 million dollars. That's way more than GiveWell can influence today and even governments are providing large amounts of aid, but they're not closing that final gap. It seems as though that's an amount that in the next 10 years you could imagine being totally closed, either through charity or general development.

Natalie: Yeah.

Speaker 9: Yeah, [inaudible 00:39:02]

Natalie: Yep. We have definitely tried to look at some of the programs that you're talking about. For example, Water Infrastructure is a program that we've look at in depth. We found that there's not strong evidence for the interventions that charities are running. I think the longer term potential solutions tend to be more speculative. They're going to not be the same evidence of effectiveness as those types of organizations. There's not as far as I'm aware a malaria intervention that is known to have a long term effect on a community in wiping out malaria. That's actually been tried there. I think in the '50s or '60s there was a large-scale malaria eradication attempt. This was ultimately not successful.

I think we're definitely open and would love to see a long-term impact of the interventions that we are looking at. I think there is long-term impact or potential long-term impact of the interventions. Just maybe not in say, disease eradication but in terms of the effects on individuals. I think long-term systems change is great, but we just haven't seen the evidence that can really be achieved by donors in our position.

Speaker 10: Really but less about long-term system change, but just about intervention [inaudible 00:41:23] long-term in general. Here's what you have experience in terms of charities such as those who are syndicates of schools [inaudible 00:41:32] child labor, which has a long payout time. I'm curious if you guys have looked into that. Also if there's any opportunity in influencing any organizations who are doing it now. Starts [inaudible 00:41:46] in 5 years you might have something to look at.

Natalie: It's definitely more difficult to collect that type of data. I think systematically there's more evidence of effectiveness for shorter term interventions. It's easier to do the study to get funding for the study. I think there's probably somewhat of a bias towards shorter term things. The one example that we've seen have a real long-term effect is in deworming. That is really the backbone of the case for deworming. We'd love to see more evidence for that. For an intervention like deworming, we expect the funding gap is probably going to be closed before you can redo the study.. we'd be interested in types of studies you could do that were more long-term, but having come across a really

good case for something that is very promising such that it's worth investing the many millions of dollars it would take in the study and the time.

Elie: In terms of where Givewell is coming from to some extent is, and as you probably know there's a lot of controversy over how well international aid really works. There's the critics. There's the supporters. One of the questions that I think is reasonable to ask is how effective could we really be at trying to solve these very thorny, hard to address problems in societies that are extremely different from our own, are thousands of miles away. The attitude that the Givewell project takes is one that says we really don't know very much, so what we're going to focus on is the set of opportunities that we can get the most information about possible in the charitable world. That definitely means that we're leaving some potential high risk opportunities on the table. That's not because we don't think they're worth undertaking. It's just not what Givewell is primarily focused on.

I personally think that what we have seen from the track record of those activities in developing countries in the last 10 years is not very strong. I think it's a reasonable bet to make on the more measurable shorter term stuff. The Open Philanthropy Project, the work that Chloe's going to talk about in a few minutes. Those are areas that are certainly not evidence-based in the way that Givewell top charities are and is our attempt to venture more into the systematic change or the higher risk, higher reward, broader impact than the direct interventions that Givewell's top charities undertake.

I think now's a good point to transition, because I want to pause and go through a few minutes of just talking about Givewell's future research plans. Then we can take another round of questions on Givewell before moving on to Open Philanthropy. Just a general organizational update at Givewell is, we've grown a lot over the past few years. We actually used to work in this building, in green spaces, in the co-working space downstairs. We were 5 people when we were working here in 2012. Now we're 31 staff members. We've really grown substantially over the last few years. The goal of that growth has been to both expand the capacity of Givewell's research. We're able to look into causes that we haven't looked into before. Also, expand the capacity for the Open Philanthropy Project, which I'll tell you a little bit about in a few minutes.

Then finally to try and get started on work that we've really deprioritized for a long time, which is reaching out to donors and telling them about the work that we do. For almost our entire history, Givewell has focused on research and deprioritized outreach, because we felt that the most important thing for us was to identify the best giving opportunities possible. We've been really fortunate that people have found us and spread the word, so we didn't have to focus as much on outreach. In the past year, we've made our first directed outreach hires.

We're now actively trying to get the word out more directly than we did in the past. I think one of the biggest challenges that we've had with all the growth that we've done is while we have grown from 5 people to 31 people, it's not like we've increased our overall output by a factor of 6. We have spent a lot of time building capacity, bringing that capacity online. I think that we're not at the point in 2016 where we're starting to turn the corner from capacity building and hiring, evaluating, training, managing to

actually having that capacity come online and produce what we hope it can. That's certainly both across GiveWell and Open Philanthropy. I think that probably over the past few years we've done maybe twice as much as we were able to do when we were 5 people, but certainly growing itself has been a huge priority of ours and certainly mine personally over the last couple of years.

Just to get more into the specifics, there's really two parts. Well, Open Philanthropy is a huge part of our future research agenda, but I want to leave that for what we're going to talk about in a few minutes. On the GiveWell side of things, there's really two big areas in which we hope to grow in the next year. One area is essentially just doing more of the same type of work that we've done but expanding it. That means following the groups that we currently recommend but doing so more intensively. Maybe doing more site visits, more audits, trying to answer question that we have as opposed to 90 percent of the questions that we have.

I mentioned this to a couple of our top charities at the end of last year, and it scared them that we had more questions that we hadn't asked them. So be it. Another thing we'd like to do is look at more programs that are within the same framework of programs we've evaluated before but just ones we have not gotten to yet due to lack of capacity. Natalie talked about deworming. This is treating these parasitic worm infections. There are other diseases in the class that's called neglected tropical diseases that we haven't evaluated but seem as if they could potentially be similar to deworming in cost effectiveness.

There are additional micro-nutrient programs, immunization programs that we've looked into to some extent but have yet to fully evaluate to the extent that we would like to. Then there are also areas that we looked into in our very early days and haven't revisited just due to lack of capacity. If you go on our website, you'll see a literature review of international education programs. We put this literature review together in the summer of 2009. We really haven't revisited it since because we haven't been able to.

In the last 6 and a half years there's been significantly more rigorous research on the program generally, but I would guess education programs. I think we now have the capacity to come back to programs like that. Education is one area I'd like to come back to. Family planning. Providing people control how many children they have is an area that I think will be challenging, because our impression is the evidence base is not particularly strong, but it's also an area that seems potentially very impactful.

Finally, surgery organizations. I think surgery charities are maybe some of the more salient groups to people in the U.S. Often they make claims of impact that are extreme. They'll say things like, "50 dollars per blindness averted." Often what they're talking about is performing a 20 minute cataract surgery that repairs a cataract in a developing country. These are groups that we've really struggled to look into before. In particular, a question we have is around the criterion of what we call room for more funding. Which is what impact marginal dollars will have. Years ago, we had this conversation with a surgery organization, where they told us how they funded a surgeon in India, who was, they were telling us, essentially worked around the clock.

He worked 18-hour days and hardly slept. When we asked them what additional dollars would do, they said, "Well, it would allow us to get the word out more or build our internal capacity," but it wouldn't enable them to perform more surgeries, because money wasn't the bottleneck. Surgeons who could perform the surgery were. That's definitely not to say that's not the last word on surgery. In some ways, it's the beginning of our investigation. Because we have to answer that question, and I think it will be working with smaller organizations in different areas, we expect this to be a more challenging area to look into, but it's one that we've already begun some work on and we hope to make progress on this year.

Then there's a whole other other area of GiveWell work that we're excited about. What we call this, for lack of a better term, GiveWell Experimental. The idea is supporting the development of future GiveWell top charities. The idea here is, what GiveWell is looking for ideally is a group that has a strong independent evidence base. A group that's already been operating for a while and has a long track record and also has monitoring evaluation data available to demonstrate that its programs are working.

You can see how what we are looking for requires other funders to have participated along the line to put us in a position where there's a group that we can potentially evaluate and recommend. With this area of our work, we're really thinking about providing support to organizations at any point along that spectrum. It could mean funding research that academics are undertaking to evaluate the evidence of a program. We've provided some funding to researchers from MIT in a program that they're not evaluating again in both India and Pakistan to see whether giving parents incentives when they bring their children to immunization camps, so basically people will periodically come to a central area to get the routine child immunizations.

The questions is whether this small incentive increases uptake of the vaccine by enough that it's worth the money you put in as an incentive. There's been one trial of this program. We recommended that Good Ventures, the foundation we work closely with, provide additional funding. Those trials are ongoing. There's a writeup on our website about it if you're interested. The long-term hope is that this helps build the evidence base for another priority program and potentially a recommended organization.

Another think along the line is at the point of an organization that's trying to get started. We've provided some funding to a group called New Incentives, which provides conditional cash transfers to encourage pregnant women who are HIV positive to take the AIDS treatment drugs that prevent mother to child transmission of HIV. They give a cash incentive at various points throughout pregnancy then after birth. This conditional cash transfer is both intended to directly transfer money to very poor people but also hopefully reduce cases of HIV. This conditional cash transfer is an area that has a significant evidence base.

We felt that this organization was very transparent, was planning to collect data in the way that we hoped. We and another funder provided some essentially seed funding. They're operating now and one stayed in Nigeria. You can read about them on our website. Then finally, we're interested in providing monitoring and evaluation support to organizations that are potentially running really great programming but just don't

have the monetary capacity or the knowledge about how to implement monitoring effectively.

This is an area that we've got a lot of questions about over the years. Do you provide consulting to charities to help them undertake better monitoring? This is really a little bit outside the scope of what I think we are able to do. We've talked a lot about capacity so far tonight. I don't think we'd be able to do it, but what we have done is find another organization whose mission is to help organizations produce high quality monitoring. This is a group called ID Insight. We provided some funding to them to try to conduct some independent monitoring of a program that's being implemented by SCI, one of our current top charities.

The reason we started with SCI is, as Natalie said, they're a group that we've had some of the least information about amongst our top charities, and so we really wanted to learn more about the effectiveness of their programming. We also wanted to find an opportunity to work with ID Insight, where because of the funding that GiveWell donors had provided, we had a little more leverage in encouraging the charity to participate. The goal here is that we see this essentially as a pilot project. If it goes well, then we'd be really excited about the potential of supporting ID Insight as a partner organization that could work with charities to help them collect the data or produce the monitoring that we're looking for.

Our big picture hopes with both the broadening of our work that's within the same scope of what we do but also the GiveWell experimental, is to have additional top charities in future years. We're largely trying to prioritize based on the groups that could be significantly better than the charities we currently have, but we're also cognizant of the fact that the gaps that we currently see for bed nets and deworming might be filled. There's always a possibility with such a small amount of groups that problems arise that cause us to not recommend them anymore or the charities themselves may run out of room for more funding with the attention donors are paying to them.

We feel it's very high priority to keep looking for identifying new outstanding opportunities. I want to again pause and open the floor up for questions either on the new research or on anything about GiveWell in general.

Speaker 11: [inaudible 00:57:02] incredibly small things that [inaudible 00:57:07] obviously is in some ways amazing from an efficiency point of view, but [inaudible 00:57:14] what's your take on that? [inaudible 00:57:18]

Elie: This is the Against Malaria Foundation. They are our current number 1 ranked giving opportunity. They're, like Collins said, a very small staff. They're basically 2 full-time people.

Natalie: They're very small. I think GiveWell's perspective is that they might be able to benefit from some additional capacity, but it's really hard as outsiders to make that decision and to know what's best for them. Largely, we trust them on them and judge them on their results. They have considered hiring additional staff and may do so in the future.



They've said to us that they believe they could benefit from additional staff, and it's really a matter of finding a good fit for such a smaller organization.

Elie: We do see it. I agree with what Natalie said. We also see a difference between AMF say, and Give Directly. I think AMF is a small organization that, at its current size, is successfully able to take donor money and turn it into bed nets in the field, but I think that's more or less the extent of what they can do. Give Directly is an organization that is able to do that for their program, take money and get it to very poor people but is also trying to build up a large organization that could potentially have even broader impacts. Some of the types of things Give Directly is trying to do, and Natalie mentioned this before, but encourage government aid programs, which, I think are notoriously poorly studied. Many people think ineffective. To at least assess whether their programs are as effective or more effective than the cash they're spending on them.

One of the outgrowths of Give Directly has been this for-profit company that they launched a year and a half ago called [Segovia 00:59:09], where, because of what they saw working with governments, the people who started Give Directly independently are starting this company that's trying to sell software to developing country governments that will reduce leakage, meaning money that gets sucked out of a system by the people in the system to increase the amount of money that's going there. This is called Segovia, the name of the company. This is the type of upside you could imagine from Give Directly that doesn't really seem that likely with AMF, given how they're operating.

Speaker 11: [inaudible 00:59:51] earlier, but there was a lot of controversy about deworming over the last year. Can you maybe summarize what? You wrote this [inaudible 01:00:01] research is?

Elie: Sure. What you're talking is something that is known as worm wars. It was one of the most exciting things to happen in global health in the last few years. Basically, the big picture with deworming is that I think this particular hubbub was around researchers who re-analyzed the data for one of the key studies for deworming. It reduced the impacts of the studies relative to what had previously been there. There was a whole brouhaha, because the researchers had previously picked one approach to analyzing the data and the followup picked another approach. A lot of people thought that the followup approach was potentially worse than the initial approach.

Suffice to say, I think that what really matters about deworming and what's important to know is that we see deworming as a high expected value program, not a high likelihood of impact program. What I mean by that is the consensus, I think if you ask most GiveWell staff members, or I'll tell you my view. The consensus is that it is more likely than not that deworming programs do not have a significant effect. When I give to SCI, which is one of our deworming charities, my expectation is that we're treating a worm, it's doing something that's good, but the effect is not very large.

But there is some chance based on some limited evidence that the program is having a very large effect, and it's so cheap that we think this program is one of the best giving opportunities. A lot of the discussion around worm wars, focused on the fact that the evidence is for a lack of a better word, pretty sketchy. It's just very limited. This is a fact

that we have always accepted. If you look at our cost effectiveness model, which is just a big Excel file on our website, the aggregate discount that GiveWell's staff applied to the headlines effects in the biggest trials is reducing the effect by a factor of 98 percent.

When we tell you that it's very effective, we are assuming that is far less effective than what is found in the trial. If you believe the trial effects, you'd end up with something that's 50 times, 400 times as cost effective as bed nets. I guess I'm going into this at some length, because I think one of the things we discussed internally this year is we think this is probably one of the things that donors following our research don't know about our view, which is that if what you're looking for is high certainty, I think Give Directly is by far the best thing you can give to.

Because it is less cost effective, but I think you are almost certainly significantly helping someone. Then bed nets are very cost effective, but there still are ways that bed net programs can go wrong. I think the single biggest concern we have right now is around insecticide resistance. What that means is that nets are treated with insecticide and they function both as a physical barrier to prevent mosquitoes from biting people, but they also are a control mechanism, where they kill mosquitoes as they land on the nets.

This part of the impact may be diminishing. The global malaria community is very worried about insecticide resistance. We've written a lot about it on our website. We've reduced our estimated cost effectiveness by a factor because of the evidence of insect resistance. It's still not enough to fully change our mind. It's definitely a way in which bed nets could go wrong. Then finally there's deworming, which I think is more likely than not, not having a big impact but is extremely cheap and maybe having a very large impact, and that makes it a pretty good giving opportunity.

Speaker 12: I was walking on the [inaudible 01:04:18] just a little north of here on Sunday. That's a store that just sold avocados. My friend and I were wondering how they can make a go of that. It seems like it's much more efficient since they have the space already rented to offer a wider range of products. If you have the capacity or is it something you were considering doing that you can get the deworming organization to work directly with the malaria organization? Or find an organization that does both? Because then you put 1 percent of the field and they accomplish essentially twice as much. Or is a third factor [3 times 01:04:51]?

Elie: Yeah. It's a good question. I think this is something that often ends up seeming better in theory than in practice. Just as an example with bed net programs, something like 3 quarters of the costs are the bed net itself. The commodity is already a pretty big part of the cost structure of the program. I don't think the savings are not quite as large as we might have thought. Then I think another issue is that different programs often have different timelines for distribution. Bed nets are trying to be distributed now every 3 years in many places. Deworming is happening either every 6 months or every year.

They're often targeting different groups of people. My guess is that for very different interventions the efficiency gained is not as big a deal as we might think. There certainly are attempts in the neglected tropical disease community, where the diseases often effect the same populations. The treatment given is very similar. It's a pill given every 6

months or every year. The target group, the population is the same to try and do what's called an integrated neglected tropical disease control. I think that's one of the possible ways. It's a possible type of organization I could imagine us recommending in the future, where they just are adding on an additional pillar or two to existing distribution.

Speaker 12: How about Give Directly setting up entrepreneurs offer a whole basket of beneficial things? You create a business for somebody, reduce the amount of money that donors have to give, because their GiveWell money is going to. The people can then purchase in their community the bed nets. You don't have to buy the bed nets. You're setting up businesses and starting an infrastructure for an economy.

Natalie: I think that's an interesting idea, and I think Give Directly has at least thought about that. I'm not sure if they actually started experimenting with providing at least information to recipients on beneficial ways that they could use their transfers. I don't think an organization that's competency is in giving out cash transfers would necessarily be effective at setting up local businesses, and that would likely increase their costs significantly. I think there's potential weaknesses in the plan.

Also just to add on what Elie was saying a little. Even within neglected tropical diseases, they use different distribution strategies through schools versus through communities. They also, if you give too many of the pills at once, they can have side effects. Even within fairly similar programs there can be a downsides to doing them together. In practicality, I think the main thing they've run into is one program being ready to go. All the pills in the country and some bottleneck happening with the other pill where it didn't quite get into place. Everything getting delayed because one piece of the program.

The more you add complexity, I think there are real downsides. I do think that the information given to Give Directly recipients could be really beneficial to them, and I'm excited to see that research when it becomes available.

Speaker 12: Thank you.

Speaker 13: Wondering if [inaudible 01:08:34] I wonder if you think that by concept including them in the top charities [inaudible 01:08:43] maybe you're saying well, they're [inaudible 01:08:51] getting this large source of funding, so we don't have to make these internal organizational changes, but on the other hand maybe that's a strength because they can focus on the things they care about, which his data analysis. [inaudible 01:09:03] see that as a strength?

Elie: Yeah, we've definitely talked about the fact that by giving AMF significant support, we are enabling them to not go out and fundraise for themselves as an example. We're excited about the fact that Give Directly is trying to build out a fundraising operation to raise money from donors who would not have otherwise given to GiveWell and its top charities. It is definitely a side effect of our support of AMF that they are much less likely to do that than they would be if we weren't supporting them.

Our general attitude is GiveWell is looking for the best giving opportunities we can find. We will tell the groups what we think about them and where we think they can improve but then, with that in mind, support them to keep doing more of what they're doing rather than try to take a heavy hand and force them to go down a path that seems better to us.

Natalie: One thing we did do this year is to start explicitly putting weight on capacity building and telling organizations that if you can use funding to expand in such a way that you change your trajectory as an organization. We think that's really high value, and we want to make sure that funding gap is filled. That resulted in us providing additional funding to Give Directly even though we thought they were a less cost effective opportunity. It's a way that we support that. Perhaps not exactly incentive [inaudible 01:10:41] but in something we started this year. We'll see how that might change things for our organization.

Elie: Do you have a question all the way in the back?

Speaker 14: Yeah, how do [inaudible 01:10:51]

Natalie: We have not looked back at Village Reach recently. I think it could be an interesting project, but probably not competitive with the other things we want to do. Vaccination in general I think is a really highly cost effective program. If there were specific opportunities for us to fund more vaccination, we would be very excited about that. The opportunities that we have looked into, the ones that we believe might have room for more funding, we've generally found that they don't have room for more funding, or the organization there's some issue with working with the organization.

For example, we have looked into [inaudible 01:11:37] and needle tetanus, working with Unicef. We had to put that investigation on hold, because they don't have the capacity right now to work with us. I think Village Reach in specific, just to go back to that, they have moved away from the specific program that we saw strong evidence for. They're doing a lot of different activities now. It's not the same. There's not the same evidence of effectiveness for what they're doing currently.

Elie: I am personally pretty interested in going back to Village Reach and just seeing what's happening. We spent a lot of time with them over the course of a couple years. It has not seemed as if trying to understand where they're currently at is the best use of our time given the capacity constraints, but it's definitely something that I could imagine us getting to in the next year or two now that we have more people at GiveWell. I think we really have historically learned a lot from going back to things that we thought in the past and just seeing how they turned out a few years later. It seems plausible that we've learned something from that.

We also are thinking about reengaging with PSI. PSI was actually our first top-rated charity. They're an interesting case, because they're a huge organization, hundreds of millions of dollars a year. They run many programs worldwide. I think the challenge for us there is going to be finding a program that has significant evidence where we really

feel like if we were to direct funding to it, more of that program would happen. Which is not always the case, because many organizations have enough unrestricted funding that no matter what funding you provide they can reallocate internally.

PSI is another organization, our first top charity that we may end up coming back to this year. Tom, last question, and we'll take a quick break.

Tom: Given the initial new [breadth 01:13:40] of research [inaudible 01:13:43] what do you envision [inaudible 01:13:44]

Elie: Yeah, I think it will probably be very similar to what you've seen. I think a possible outcome of looking into education, family planning and surgery, which, to be clear, I would not guess we will finish all 3 of those next year. I'm hopeful we'll finish one additional thing like that. I think the possible outcome is 3 long research pages that talk about all the questions we have and the charities we looked at and the fact that we couldn't answer questions. Get some questions answered. A few charity reviews that describe our process and the questions we asked and the information we got, but ultimately end in an unsatisfying way, where we don't have compelling answers. If you go on our website now, you'll find a lot of that that goes beyond the 4 top charities we have.

It's just hidden in the guts of the website. It's all the work that stands behind the 4 top charities. I definitely think there's a point in GiveWell's future that we will actually get better at presenting our material, the interesting research we've done, to the outside world in a way that doesn't require people reading the equivalent of a 50-page research report to know what we think. We see this as one of our most important outreach priorities but also probably one of the hardest to undertake effectively.

Because getting the right, pulling the most important pieces of a research project and presenting it in a way that we're comfortable frankly given our brand that is both trying to present what we know but also be open about our uncertainty has been a challenge for us. We call this project the media level depth content for GiveWell work. I hope that we get there, but I'm not optimistic that it's happening in the next 12 months. Let's use this as an opportunity to take a few minute break. Feel free to grab more pizza if you want and grab some more drinks. We'll reconvene in a few minutes.