

GiveWell NYC Research Event June 26, 2018 – 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: Thank you, everyone, for being here for the GiveWell Research event tonight. This is a really awesome turn out and awesome crowd and it's really nice to see some old faces that I've been talking to for a while at these research events as well as some new faces. So welcome, everyone. I'm Catherine Hollander, a research analyst focused on outreach at GiveWell. So, just quickly for the agenda this evening, I'm going to speak about the latest research that GiveWell has been working on in 2018 and then leave plenty of time for questions during this session and then we'll break. There's still some pizza left, so hopefully folks will help themselves to more pizza. And mingling and hanging out and speaking to me and the other GiveWell staff that are here is strongly encouraged. We want to meet everyone and chat with you if you have questions and get to know why you're here and how you found us. And I'll introduce the other GiveWell staff at the end. We're all excited to meet you. And that's sort of the format for the night. We have this space 'til 9 o'clock. This part of the talk will go for about an hour and we'll chat and hang out after that.

01:16 CH: As a reminder for those of you that have been here before and for those of you that are new, we generally record GiveWell Research events to share on our website with folks who couldn't be here today. So I'm recording currently. If you ask a question during the Q&A period that you would prefer not be included in the recording or transcript that we produce, please just email info, info@givewell.org in the next two days and we'll remove it before we publish it. But we like to be able to share for folks who are not based in the New York or San Francisco area, what we talk about at these events. I'll also repeat your questions for the sake of the recording. So I'm going to start just with a very brief overview of what GiveWell is and what we do. If some of you are familiar, others, you are here for the first time, so just to make sure everyone has kind of a basic overview of what GiveWell is working on. If you have any questions just about the basics of how GiveWell operates, I'm very happy to chat afterward or feel free to ask in the Q&A. All questions are welcomed.

02:25 CH: And then hopefully that will set some of the context for the research that we're working on this year, so I will then talk about some of the major research questions that we have on our mind related to the current group of Top Charities that we recommend, that we're focused on answering this year, some of the work that we're doing to identify new Top Charities and also some of the work that we're doing to incubate and develop potential future Top Charities. So just starting quickly with the kind of overview of what GiveWell is. Our focus is on finding charities that have the highest impact per dollar and recommending them to anyone who wants to use our research. So the product that we produce is a shortlist of top recommended charities, that's on our website and available to everyone.

03:12 CH: These Top Charities work in the poorest parts of the world. They work in global health and development. And the types of programs that they work on are things like treating intestinal parasites, distributing insecticide-treated nets to prevent malaria and distributing direct cash transfers. We currently have nine charities on our list. These charities are selected based on questions that we try to answer to understand who to recommend and how to rank them and there are three basic criteria that we're looking at when we're thinking about which charities we might recommend. The first is, what's the evidence of the charity's effectiveness? And we actually start that research at the program level. So a program is anything, is the thing that the charity is doing but independent of what the charity itself is. So just to make it less abstract, a program is distributing insecticide-treated nets. The charity that GiveWell recommends that does that is the Against Malaria Foundation. So we first look at the evidence for the program, whether if something like distributing insecticide-treated nets is effective. The second thing that we look at is, how well it works.

04:20 CH: So we're interested in how effectively the charity is able to deliver that intervention, how strong their ability is to demonstrate that they have reached the people that they are intending to reach and what we know about those outcomes. That's our first criteria. The second that we look at is the cost-effectiveness. So, as I mentioned, we're focused on trying to find charities that have a very high impact per dollar and so we are extremely interested in what are some of the cheapest ways that we can find to help people. And so we spend a lot of time looking for the programs that we believe can do the most good per dollar donated. So cost-effectiveness is a really an important criteria of ours for all the charities on our list. The final criteria that we look for is how much additional funding a charity needs and whether they can effectively put additional dollars to use. Since the product that we're producing is a list of Top Charities that people will donate to and direct dollars to, we want to make sure that every additional dollar can go a long way. And so, as I mentioned, we use this criteria to identify our list of Top Charities.

05:26 CH: The nine charities on our list currently work on a variety of programs and we're interested in recommending

additional charities. We find additional charities that work on other programs that meet our criteria. And in 2017 we estimate that we directed around \$117 million to the Top Charities that we recommend. So that is something we're very excited about and I think it's important to many people in this room, we really appreciate that and are hoping to continue to move a lot of money to the great groups that we recommend. And so I mentioned that the first thing I wanted to focus on is to talk about the biggest research questions on our mind in 2018, the things that might change our opinion the most about the current group of Top Charities that are on our list. And so, there are three major questions that we're looking at right now.

06:20 CH: The first thing that we're looking at is research that's been done on the impacts of cash transfers. So, cash transfers are the intervention that is implemented by one of our Top Charities called GiveDirectly. They've been on our list for a long time, and what they do is that they deliver cash directly to some of the poorest people in Kenya, Uganda and Rwanda. And this year, some new research came out on the impacts of cash transfers that had two major suggestions. And before I get into those, I just want to start with our bottom line, which is that we are planning to look closely at this new research, and that we expect that our views could change a lot as a result of it. But our best guess is that we will continue to recommend GiveDirectly once we're done with this review of this new evidence.

07:14 CH: So, the research that came out basically had two... It suggested that cash transfers might be less effective than we previously believed, in two different ways. The first is that cash transfers might have negative effects on the people who do not receive them, or spillover effects. We currently assume that cash transfers do not have a major negative or a positive spillover effect. So, if your neighbor gets a cash transfer, we don't assume that you're much better off and we don't assume that you're much worse off. This study that came out this year, and if you follow development econ blogs, you've probably heard about it, it caused a bit of a dialogue to happen there. It was a three-year follow-up study of a GiveDirectly program that distributed cash transfers of about \$300 to about half of the people who live in a given village, and this study suggested that the people who did not receive the transfers were worse off.

08:15 CH: And so for GiveWell, as a recommender of cash transfers, we are definitely very interested in the discussion. If it turns out that cash transfers have these negative spillovers, that might really influence our bottom line. But there are a few reasons that we think that this particular study isn't directly applicable to the work that GiveDirectly is doing today. And we're planning to wait for the results of another major study that's being done on a version of GiveDirectly's program that is much more similar to what they're doing now, and we expect results from that study to come out in the next few months. So, the version of GiveDirectly's program that they're currently implementing in Kenya and Uganda is to give \$1,000 transfers to all members of a village. So, I mentioned that in the study that came out this spring, they're dividing the village and only giving transfers to about half of eligible recipients within the village.

09:10 CH: GiveDirectly is now operating in a model where they are making transfers that are both larger, about \$1,000, and to the whole village. So, negative spillovers may still occur. Our best guess is that if they do, it would be between villages rather than within villages in the current model. But we're waiting on this major study that I mentioned is coming out. We expect, I think, mid-line results from the study in the next few months to do our really big deep dive into what we believe about the effects of spillovers. But this is something that we're watching quite closely and that we plan to do a deep dive on before the end of the year.

09:45 CH: And the second... I mentioned there were two questions that got raised about cash transfers due to new research this year. This second question was on the duration of benefits. So right now, we think that people benefit from receiving cash transfers in two different ways. The first way is that you benefit from being able to buy things right away. So, things like food. Immediate consumption benefits. And the second way is that you are able to invest some of your transfer and have benefits that might accrue to you 1 to 20 years in the future. And the new research that came out this year suggested that the medium-term benefits might not be as significant as GiveWell assumes in our current model.

10:28 CH: We have not done a deep dive on this yet either, as we're planning to revisit it as part of our big look at cash transfers once we have the new results of the midline study that I mentioned that's coming out looking at the program that GiveDirectly is implementing today, but we think that this is a really important question. With the duration of benefits and how long we expect benefits to last, we think that the floor for how low our cost-effectiveness could fall is about 1.5 to 2 times, and the reason for that is because we currently assume that about 60% of transfers are just consumed right away. So, even if you completely removed the 40% benefit that happens over the medium term, we still end up feeling pretty positive about the benefits of cash transfers.

11:14 CH: But certainly this is something that could make GiveDirectly look a bit worse. And so, we're planning to

follow this pretty closely this year. And so, right now, we're at the sort of medium-term update on this, but we don't have our main conclusions that are planning to publish that by the end of the year. And so, the major bottom line is that this might change our view significantly. We could imagine that will happen in a positive way. There is an argument to be made that there might be positive spillovers for cash transfers or a negative way, and we're not yet sure. But our current best guess is that we will continue to recommend GiveDirectly.

11:48 CH: So, that's the first major research question on our mind, which is on cash transfers. The second is related to insecticide resistance. So, I mentioned earlier that we have a Top Charity called the Against Malaria Foundation that distributes insecticide-treated nets to prevent malaria. And these nets work in two ways. The first is that they provide a physical barrier, so they block mosquitoes from biting you, and the second is that they kill the mosquitoes because they are treated with insecticide. And we are concerned about a potential increase in insecticide resistance. Our impression is that it's increasing across Sub-Saharan Africa, and we've done several reviews over the years of studies that have been done to try to better understand how potentially significant this problem could be. However, what we found in those reviews is that there aren't any high quality studies that estimate the impact of insecticide resistance on insecticide-treated nets' effectiveness in the field, and so we don't currently have a lot of information about the scope of this problem, or this potential problem.

12:58 CH: Because the Against Malaria Foundation is one of our Top Charities, we think that this is a really important question for us to be asking. And this year something that's new, is that a new type of net is being tested, it's a new type of insecticide-treated net called a PBO net. The way that this works is that, and I'm not a biologist, so please forgive the very simple explanation. But the way that I understand it is that a PBO is a synergist that works by blocking the enzymes in the mosquito that sequester the insecticide and prevent it from having the toxic effects on the mosquito. We've seen some initial evidence that looks promising on the fact that these nets might work well, and there's been one randomized control trial, which is a particular type of study that GiveWell uses a lot to look at the impacts of the health interventions that we recommend, where basically you divide a group of people in half, a randomly assigned half to get the treatment, which is in this space would be a PBO net, and half to the control, which would be people receiving the standard insecticide-treated net, and then look at the differences between those two groups.

14:07 CH: And so, this initial evidence that we've seen, we think was positive, and we also think of this as being pretty interesting as a potential proxy for the current level of insecticide resistance, since we can compare how effective these new nets are with the PBO to the old insecticide-treated the nets, and we actually, it's not considered ethical to randomize people to have untreated nets, and so we don't have that information about how current insecticide-treated nets compare to untreated nets, or PBO nets to untreated nets. That's why we think this is an interesting proxy for that.

14:46 CH: And our current kind of next steps is we're waiting for another randomized controlled trial of PBO nets in Uganda to be completed. The Against Malaria Foundation is distributing those nets, that's the Top Charity that we have that I mentioned, and we're planning to then revisit our estimate of the relative effectiveness of PBO nets to standard nets and to incorporate that into our understanding of the efficacy of nets, and also to incorporate into our cost-effectiveness model the areas in which the Against Malaria Foundation is distributing these nets as opposed to the old nets, since we expect that there would be differences both in the cost of the nets as well as the impact of the nets and that would factor into it kind of our bottom line that I mentioned that we produce, which is a look at how expensive it is to produce a particular good outcome, in this case, averting deaths from malaria.

15:47 CH: The third question that I wanted to talk about, is one that we think is really challenging, which is that donors, anyone who's giving to charity and anyone who, like GiveWell's in a position of making recommendations or suggestions for where someone might give to charity, has to make a challenging moral judgment call between different positive outcomes to try to come to some sort of bottom line to say that you should donate here to have the most impact. And concretely what I mean is that we currently look at charities that work on a very wide reach of interventions.

16:29 CH: This is from the Against Malaria Foundation, where we think that the benefit is averting deaths from malaria, to cash transfers, where we think the benefit is having the ability to buy more things, to groups that work on deworming, which we think increases incomes later in life, so that there's a very broad range of outcomes that we're looking at, and we're trying to come to a bottom line and say, "Here's where we think your dollar can do the most good, can deliver [16:54] ____." We have to make a trade-off between those outcomes to say, how valuable do we think it is to increase someone's income relative to something like improving their health, and it's really important for us to get good inputs here on how to make those decisions and describing funding decisions and recommendations that we're making.

17:14 CH: And right now, what we've found is that the information that exists that we can use to inform this decision is quite poor. There's not a lot of information out there about how people in the parts of the world where GiveWell Top Charities operate would make these trade-offs. The current way that we approach this question is sort of the best that we can. We have staff members input their individual inputs into the model that we use to look at the cost-effectiveness of different interventions, and those inputs are shaped by things like global health research, philosophical world views, thought experiments and research. Each member of GiveWell's research team might trade those values off differently based on these factors, or other factors that they consider important.

18:04 CH: But we think that this is sub-optimal, since it does not include information on the preferences of people who receive the benefits of our Top Charities. And so, we first looked to see whether there was research that was done on this and found that that there wasn't. But there is some research studies on how people in rich countries make these trade-offs, but very limited information on how people in the poorest parts of the world would. And so in the medium term our plan is to go out and gather this information ourselves. And so we're working with this group called IDinsight, which is a group that we've worked with for a few years through our Incubation Grants program, which I'll talk about in a minute, and they're a unique group that's focused on providing decision-relevant information to policy makers and implementers, and so we're working with IDinsight to go out and survey beneficiaries of our Top Charities to see how they would make those trade-offs between income and between health.

19:02 CH: And so they've done this year a pilot, which we have some initial results from, that they conducted in Kenya earlier this year, I think it was in February, where basically what we learned from this pilot is that it's very challenging to ask these questions. It's very hard to design a survey to give you information that is easy to understand and interpret. And so, for example, one of the questions that this research focused on was this question I've been talking about, the value of averting deaths versus increasing income and consumption. And it was very unclear what to make of the results, because IDinsight used a few different framings to ask this question to try to understand whether people might answer the question differently if it's phrased differently and framed differently.

19:48 CH: And they got very different answers from those different framings, so it's hard to know what we should make of what we got back from that survey, and so we're considering now with this information, having done the pilot, funding another pilot that would use different approaches and framings, so that we can see whether we might be able to learn more from that that would affect how we think about these trade-offs in our model. We're also interested in seeing whether a different population, maybe in a second pilot, would give different results. So if the first pilot was done in a part of Kenya, maybe people in a different region would have different answers. It was also primarily... The first pilot was primarily speaking with women because they were the easiest to survey. And so, we might also wonder whether men would give different answers, and so we're currently in the process of thinking through how we might design a second pilot to help us get better information to inform this really important piece of our cost-effectiveness model, where we think we just currently don't have very good information and are doing the best we can with what we have but are excited that we might be able to fund and generate more information to inform this really important outlet.

21:02 CH: The final thing that I wanted to mention about our current group of Top Charities, that I just think is really important to kind of keep in mind as sort of a big picture check-in, sort of where we are in the world and how things are going, which is that one of the projects we did this year was a review project to look at the overall size of the total global need for insecticide-treated nets, and also another anti-malaria intervention called Seasonal Malaria Chemoprevention, which is another program that GiveWell recommends, where you're giving preventive anti-malarial drugs to young children during the time of year when they would be most likely to catch malaria. And so we are very interested in funding these two interventions, as we think they're both very effective ways to prevent malaria.

21:53 CH: And so, we wanted to get a better sense of what the total funding needs are in each of these phases, to see how much additional funding might be necessary to fill the global gap in these areas. And what we found was for nets, we used information from the RBM Partnership To End Malaria, which was formerly called Roll Back Malaria, and for the Seasonal Malaria Chemoprevention data, we looked at information from the Malaria Consortium, which is our Top Charity that we recommend that works on this program, and wanted to find out what they believe the funding gaps are in these spaces from 2018 to 2020. Our conclusion based on their research is that there is a gap of about \$648 million to deliver nets to all the people who need them. So \$648 million, which is significantly more than what we moved to our Top Charities last year. It's a very large gap.

22:45 CH: The largest gap there was in Nigeria, that's about half of the gap, so \$330 million, roughly, of need that we see in Nigeria. And, for Seasonal Malaria Chemoprevention, the anti-malarial drug program, the gap was \$165 million. Again, this is a lot more than the total amount of money that we directed via GiveWell, and I think just speaks to where

we are currently. We hope that we can direct more funding to these opportunities and that there still is quite a large need for what we believe to be some of the most evidence-backed and cost-effective ways of helping people. We do think that there are reasons that the gaps might shrink relative to this initial estimate that we have. One was that the estimated gap for 2017 was relatively small. The RBM Partnership believe this was due to large contributions from a funder that won't repeat. This year, this no longer is available, but that could be one reason that the gaps could be smaller than we think.

23:45 CH: Another is that Nigeria is working to get funding from development banks which could reduce their need and, as I mentioned, they're a significant portion of the need for both the nets and also for Seasonal Malaria Chemoprevention. And then the final reason, the gaps could be a different size than we initially estimate is a country's own responsibility to reallocate funding between their different programs. So that's the usual GiveWell caveat that this is the best guess that we have and the best conclusion that we have currently, and then here are the reasons why it might be different from that, but we still think the kind of big picture bottom line is that there is a significant need here and so as donors we think is really important to have the total picture in mind that there still is a long way to go. That's the kind of broad update on our current group of Top Charities. I might pause here and take a few questions on this before I move on to the work that we're doing to identify new potential Top Charities. If you have questions on the new Top Charities, we'll take those after, but let's take a quick pause here and talk about the current Top Charities.

24:50 S?: As I understand it, the concern about the insecticide resistance is that insecticide resistance could make the bed nets less effective. Is there a risk in the other direction that you would be increasing insecticide resistance by spreading the nets?

25:04 CH: So the question is, is our concern with insecticide resistance, that insecticide resistance makes the nets less effective and are we also concerned that we could be increasing insecticide resistance by using the new nets? So the answer to both of those is yes. We think that the concern with insecticide resistance is that the nets that are treated with a particular type of insecticide will be less effective. And so the idea is that there are different types of insecticide, so mosquitoes would become resistant to one type of insecticide, but not necessarily a second type of insecticide, although it's always... The resistance could develop in both directions or once we shift to using a different type of net, then you become concerned that resistance could grow to the new type of net. You're always kind of trying to out run the mosquitoes' development of insecticide resistance.

25:55 CH: And yes, any time that you use insecticide, you are increasing the risk that mosquitoes will evolve to be resistant to it, but I think that because you can introduce new types of insecticide that mosquitoes aren't yet evolved to, we see that there are significant benefits doing that over something like an untreated net or sticking... As I mentioned, we have... One study that's been done on the new nets that shows that they're more effective than the current group. And so that's what we want to see and to implement. And so we think that that's sort of the best path, but certainly PBO resistance could occur down the line as well, and we need to think about that too.

26:37 S?: Who is IDinsight and who actually works there in [26:37] _____.

26:44 CH: Yeah, so who is IDinsight and who works there? IDinsight is a very large international NGO. They have an office in San Francisco, which is where GiveWell is based, but they also have a number of country offices. I'm not positive they have a country office in Kenya, but they employ a number of field researchers who are the ones that go out in the field to conduct survey work. We see them as really unique in the global health and development space, because a lot of the research that GiveWell relies on in making our recommendations comes from groups that are very aligned with us that specifically focus on global health and development and also academia, and we think that academics aren't necessarily always incentivized to provide the information that's sort of most decision-relevant for GiveWell, either because it's not aligned with their career incentives.

27:34 CH: We're very interested in things like very targeted research on specific areas that our charities might actively consider working in. And so we've found that academia doesn't always have the kind of most decision-relevant information that we need and IDinsight is unique because they focus explicitly on providing this type of targeted information where we can come to them with a question as sort of a meta-research evaluator and that they can go out on the ground and help us answer the question. So GiveWell doesn't do on-the-ground research ourselves, but IDinsight is coming in as a very aligned group that we're working with to help fill those gaps in our knowledge and working with us on that.

28:13 S?: I'm trying to estimate the value of a life or a healthy year. Have you guys tried looking at how much of the

current giving that GiveDirectly has done has gone toward medical spending or preventative medical spending, and trying to extrapolate from that. And/or has anybody ever looked at the developed world where you see jobs that people take, riskier or less risky from mining or farming, even citing, [28:40] ____ similar incomes in India or something.

28:43 CH: Yes, the question is, have we looked at the information that GiveDirectly currently collects on whether people are spending money on health-related items, and also, have we looked at things like revealed preferences surveys in other parts of the world that look at whether people are more willing to take a particular type of job at a particular pay with a higher risk of death, which is one way that then I understand that in richer countries, this question of how people trade off income and health is often done, which is just looking at how much more do you need to pay someone to work in a job with a higher risk of death?

29:19 CH: We're not sure how applicable the research that has been done on that is to the context in which our charities work, which is the very poorest parts of the world, but we did look at that type of research when we were thinking about what information is already out there. So we have seen it, but we are interested in sort of the more context-specific information to where our charities work, as we could imagine that people might make these trade-offs differently in the poorest parts of the world relative to the rich countries in which this is done, or the research is often done.

29:50 CH: And then the first part of your question, when we looked at the sort of breakdown sending of GiveDirectly recipients, the answer is yes, and in our GiveDirectly review, you can see the research that we have on what people tend to spend their transfers on, but when we at GiveWell are modeling the benefits of receiving cash transfers, we're focused on the increasing consumption as the sort of primary outcome of interest, so we're not modeling health benefits via the potential to send them health goods kind of explicitly in our model of GiveDirectly, but we have seen the information that's available there, and that might factor into individual staff members' thinking about how they would trade those two off.

30:37 S?: Yes, so Bill Gates is trying to eliminate malaria entirely, whether that's eliminating mosquitoes or something else. Does GiveWell have a perspective on whether malaria will be eliminated, after you do your evaluation of things like the medicines and...

30:55 CH: Yes, the question is Bill Gates is spending a lot of money on malaria, does GiveWell have a perspective on whether malaria is likely to be eliminated? I think GiveWell hopes that malaria is eliminated, but we don't currently factor that into how we think about the benefits of providing someone a net today. I think as an organization we have a very strong focus on what can achieve the most good donating today and right now, malaria does kill a significant number of people every year and does have these very large funding needs, and so we think it's really important to help people today where we know that there is this evidence-backed intervention, and that if 10 years from now, we can't recommend the Against Malaria Foundation because malaria has been eradicated, we consider that to be a very successful outcome. We would be very happy about that. We are glad to take things off of our list if the need no longer exists.

31:46 S?: But for the medicines, it's a bigger and different problem than the nets. So I'm with you on the nets, but if somebody's like resistant for 20 years, you'd think like the future timeline would play into that, into the value of like a vaccine or something, if that was possible.

32:03 CH: For the medicines, are you thinking of something like, is there a preventive anti-malarial drug? So I think those drugs are administered very regularly, I think the goal is once a year, maybe twice a year. And it's the kind of thing where you have to have it re-upped regularly, so it's not a vaccine that lasts for 20 years. This is like a very near-term intervention that prevents people from getting malaria today, sort of in a similar way to... We think that insecticide-treated nets last for about two years before they degrade to the point where they are no longer as useful in preventing malaria. So nets get holes in them, that type of thing, that makes them less effective over time, and so we think that these are sort of short-term preventative measures and that anti-malarial drugs would fall into that camp. So it's not sort of a long-lasting vaccine, but rather a shorter-term intervention.

32:53 CH: And I think just looking at the time, I'm going to talk a bit about our potential Top Charities. We can definitely come back to questions on our current Top Charities post that, but just to make sure that you get the full download, we'll talk about our potential Top Charities now and so come back to these questions in a little bit.

33:13 CH: So as I mentioned, those are all of our big research questions, or some of our big research questions, related to the nine groups that we currently recommend. There's basically two parts to our research process. The first part is

scanning the world of potential interventions. So I explained earlier, an intervention is the term that we use, we use it fairly interchangeably with the term program, which just means the sort of abstracted thing that you can do to help someone. So distributing insecticide-treated nets, giving vaccines, distributing cash transfers, treating intestinal parasites, that's separate from a particular charity.

33:50 CH: And so the second part of our research process is looking at the particular charities that are implementing those programs. So step one, looking at all the potential programs that we could potentially recommend and then step two, looking at the charities that implement those programs. And the reason that step one is looking at all of the interventions at the intervention level is because we found that the differences in the evidence of effectiveness and cost-effectiveness, so two out of those three criteria that I mentioned at the beginning of the talk, those differences are primarily driven by which program a charity is implementing. So it's a much bigger difference in a charity that is implementing the program of insecticide-treated nets and a charity that is implementing a program of providing surgeries in the United States.

34:39 CH: The differences in the programs will drive a lot of the differences in the cost-effectiveness between those two groups, so that's why we think it's sort of the logical first step to start at the intervention level. So we focus on our first part of our research on which interventions have the best evidence and are the most cost-effective. And so at GiveWell what this means in our office is that we have a huge spreadsheet of... I don't know the total number of rows that are in the spreadsheet right now, some of our research analysts here might be able to help, but something like 100-plus different potential interventions that we might look into that we are prioritizing to try to think about how much time we should spend researching them, how promising they are and whether we should ultimately elevate them to the level of what we call a priority program, which is the point at which we say, we've looked at the evidence for this intervention, we think it is strong enough that we're excited to start the next phase of our research, which is finding a charity that is implementing that program.

35:35 CH: And the way that we add things to our dashboard is sort of a combination of reviewing new academic literature that's coming out. We follow the journals quite closely like the Lancet, the New England Journal of Medicine to see whether there are programs that are popping up there through new studies that we should have on our radar. But we also have conversations with charity representatives and academics that we think are helpful to us in terms of keeping an eye on what new research is coming out and what we might want to investigate further.

36:04 CH: And so, one intervention, just as an example, that's recently come out of this process is screening and treating pregnant women for syphilis infection. Our current best estimate of the scope of this problem is that there are about 1.4 million pregnant women that have syphilis, and the complications from not treating that are stillbirth, neonatal death and potentially long-term complications for the child. There is a very simple point of care, really. It can be administered right where the patient is, intervention, which is diagnosing whether a woman has syphilis and treating it with a simple dose of penicillin. There's a cheap, simple way to address this problem that we think seems like it's potentially something we want to look into, and our best understanding is that this intervention, syphilis screening and treating, is currently neglected within the global health and development community.

37:06 CH: It's a fairly small problem in the scheme of things. And if we think there's probably about \$10 million dollars worth of funding need to treat this group of women, which might explain why we think that it's fairly neglected, meaning that it's not receiving other funding, why we think that GiveWell might want to recommend it so that we could direct dollars to it. We spoke with the person at the World Health Organization who is in charge of this program, and they told us that they spend less than half of their time per year working on this, and other funders that we've talked to don't seem to be very focused on this program. So we think that this might be the type of thing that has the combination of the strong evidence that the screening and treating program works, that it might be potentially competitively cost-effective with the current Top Charities that we recommend and that it might be neglected.

37:53 CH: In other words, that there might be room to fund this program. So because of that, we are planning to continue to research this program that's kind of come through our prioritization dashboard, and one thing about this program that we think might also make it a fit for GiveWell is that we don't think that it is likely to appeal to mass market donors. And it doesn't seem like any of the existing entities that we might expect to take action on this are likely to take action. We think there really might be a gap here that we can fill. The problem that we have found with this that if we decided to fund it, we're not sure who to direct funding to.

38:28 CH: So I mentioned the two steps of our research process, intervention prioritization and then looking for charities that implement them. We have found through GiveWell's history that sometimes we run into a sort of road

block where we've identified a program that we're interested in, but not a charity that we can direct funding to that's implementing the program. This problem is not unique to syphilis, but it is something that we're interested in addressing. And so the way in which we're planning to address this and have started work on addressing this, is through our Incubation Grants Program, which I briefly mentioned earlier, is the way in which we've been working with IDinsight, that we were talking about earlier. And so we are interested in whether we as GiveWell can direct funding to things that would enable the creation and scale up of a charity that others could eventually support in the form of a GiveWell Top Charity.

39:21 CH: So can we fund organizations at an earlier stage that we would recommend them as Top Charities to donors and as we pretty strong criteria for what we look for in our Top Charities, we really want charities to be very evidence-backed and have a very good track record, but through this program can we support things at this early stage, to scale them up to the point where we might recommend them as Top Charities. And so one of the grants that we're currently thinking about in this category, related to the above syphilis treatment, is a grant to a group called Evidence Action, which you might be familiar with if you have followed GiveWell for a while. They're the parent organization of two of GiveWell's current Top Charities, the Deworm The World initiative and a group called No Lean Season, which we added to our Top Charities list last year, and is a seasonal migration program in Bangladesh that was actually an Incubation Grant recipient itself back in 2014, that kind of came through this newer program that I mentioned where we're trying to grow potential Top Charities.

40:23 CH: And so we're considering giving Evidence Action a grant for a part of their work called Evidence Action Beta. And the idea here is to go from Evidence Action's current model of picking programs that seem promising and sort of slowly scaling them up, to working with them as sort of an implementer or an incubator that has a broad portfolio where we can look with them at a lot of ideas kind of at the beginning and hopefully lead to more Top Charities over the long run. And so syphilis screening and treating is a potential program that could be part of a Evidence Action Beta grant, and overall, with Evidence Action, they are a group that we have worked with for a long time, and feel very closely aligned with, and so we're excited about the possibility of making a grant like this to help us really not only support things via recommending them to donors that we think are great, but to enable the development of things that we think could be really exciting for donors in the future, once there is more evidence there.

41:25 CH: And so that's kind of the overall [41:29] ____ program having more Top Charities and better understanding the current group of Top Charities that we recommend. And I think this type of grant is really illustrative of the work that we're doing there.

41:41 CH: So I'm going to pause here for more questions on this or the earlier part of the talk on Top Charities or any other things that are on your mind. We have about 10 more minutes. So I would love to answer any more questions.

41:57 S?: When you were talking about GiveDirectly, you said that there were some negative spillover effects. Can you talk about some of the outcomes that you're familiar with, like what you mean by that, so your neighbor non-recipients of direct dollars, how, on what outcomes are they faring poorly? Is it negative attitudes or global wealth or development issues? Do you have any... Are there any theories as to why that's happening?

42:28 CH: Yes, so the question is basically more information on the potential negative spillovers for GiveDirectly. So, we know GiveDirectly has done some research themselves on the sort of emotional spillovers, I'm not sure that that is the correct term, but thinking about are there jealousy impacts, are there impacts on those feelings. We talk about those in our GiveDirectly review, so there's more information there that I can share with you. And then on the recent study, I think that's looking primarily at economic outcomes and economic spillovers, although I'd have to double check that, so I'm happy to get back to you on what concretely they looked at. I know for us, since the consumption impact is really important, we would be really interested to know if there are major economic benefits or negative impacts on neighbors, but also if we learned that there were significant psychological well-being impacts, we would be curious to incorporate those as well. And so we're waiting to do our really deep dive on this until that larger study comes out, since we think it'll be more telling for what we should do, as it more closely resembles the current GiveDirectly program.

43:36 CH: So yeah, I don't have a lot more information off the top of my head on the shorter term results that we looked at, but I'm happy to follow up with you on the [43:42] ____.

43:48 S?: I'm wondering about the surveys that you're running about GiveWell's preferences for different outcomes and... It sounds very, very strange to me that you ask questions in a different way that are [44:03] ____ like that, that really increases the environment where it's really easy to ask questions [44:09] ____ basis, so you certainly true... And

I'm wondering about, just like framing and going and asking people what they want in this abstract way, especially for different interventions, which are like these targeted and different actual... Like you can actually go to these people and say, "Would you like these bed nets, which we believe have this effect, or would you like this cash, which we believe will have this effect." [44:34] ____ has to explicitly choose between different people receiving outcomes. So I'm wondering about, like, it seems like really [44:47] ____ do something that's evidence-based or sort of...

44:53 CH: Yeah, so the question, or the concern is, can you do something that is evidence-based and feel like you really are confident in the answers that you get, knowing how many challenges there are in framing these questions or how people might respond really differently to frames and how well do people understand the trade-offs that they're being asked to make. And really, you know, this whole project seems really challenging. How do we get anything useful out of it?

45:18 CH: Yeah. So, I mentioned they used two different framings in the initial pilot that they did that gave these very different answers, and maybe just to kind of give a more concrete sense of what they are. The framings were called the "taking framing" and the "giving framing." The taking framing was asking how much money would you pay to basically avert the death of, they asked about a number of people, but like your spouse, a child of a particular age, yourself, and the giving framing is a donor is interested in funding cash transfers, or something that will save the life of a spouse, yourself, children of different ages of yours, and looking at the different answers that you get to those two framings.

46:00 CH: A concern with the taking framing is that people might be... They might feel limited by what they believe that they can borrow in current markets for what they would answer for that amount of money, and the giving framing gave a really different answer where there was basically no amount of money, I think is what they concluded, that would be traded off against averting a death as the median answer. And I think it was about \$1,000 with the taking framing. So just getting really different answers makes this really complicated to interpret. And so I think we're aware of the challenges in doing this, and I think trying to approach this and thinking about what information could we get from this that we would believe could be a meaningful input in our cost-effectiveness analysis, and how much skepticism should we approach the answers with.

46:48 CH: So I think we're very aware that this is a really challenging question to answer and are planning to treat the responses that we get as carefully and thoughtfully as we can, knowing that these are the challenges of getting good information in a survey like this, where there is so much potential to have confusion or complications in answering.

47:10 S?: Are there any overlap populations where there would be a single person who could potentially benefit from either this program, [47:17] ____?

47:21 CH: Yeah, are there any individuals who would, say, potentially receive a cash transfer and also an insecticide-treated net? I don't know off the top of my head. I would guess the answer is yes. But I'm not sure even, I'm not quite sure how you would frame that question the way we're... Yeah, it's all very...

47:38 S?: Based on the amount of... In terms of per unit of giving...

47:46 S?: If I'm going to give somebody \$10,000 worth of nets...

[laughter]

47:51 S?: You're going to give them like, I don't know, [47:56] ____.

47:58 S?: Well, or you'd have to divide it by the appropriate number.

48:07 S?: To what extent are bed nets sort of funded by charities such as the Against Malaria Foundation versus [48:14] ____, I guess, like governments [48:15] _____. You mentioned that as sort of a funding gap. I guess like would it be effective, I guess, to [48:25] ____ governments to... Yeah, sort of spend...

48:31 CH: So the question, I'm just repeating this for the sake of the recording. The question is, how much of current net spending is done by governments versus groups like the Against Malaria Foundation, and also, have we thought about encouraging governments to spend more on nets? So, I'm not sure off the top of my head what the breakdown of government versus AMF funding is. I'm happy to follow up with that since I don't know it off the top of my head. In

terms of whether we might try to encourage governments to do more with net spending, the Against Malaria Foundation does work with some governments. They have to get government agreements to work in different countries, but in terms of political advocacy, I think we are interested in whether that type of work is the type of thing that we can assess going forward, but I think in this particular streak I'm not aware of groups that we would assess to do that type of advocacy.

49:32 CH: I think malaria is a pretty well-known issue and I think country governments do fund it to some degree, so I'm not quite sure what the advocacy space would look like within this area, but we do know that we can give money to the Against Malaria Foundation and believe that they'll scale up net use, so I think for now, we do think that is a really good option for us to provide funding for.

49:58 S?: I wanted to ask about the motivation behind doing the IDinsight work. So, is the assumption that people who are the targets of the intervention, who have benefited by interventions, have better information about how their lives would be improved by the interventions versus money? Like their lives continuing, being lengthened, or getting more money? Is that the assumption? That they have better information? Or, is the assumption that even if they had no better information we should defer to their preferences over our own preferences?

50:30 CH: So the question is basically how we would think about the responses we would get back from a beneficiary preferences survey. If we would assume that beneficiaries have the equivalent information that GiveWell has about the impacts of the different interventions or whether we would weight their preferences above GiveWell staff members.

50:51 S?: Right.

50:53 CH: Yeah. I think either that, the sort of weight to which individual staff members might put on beneficiary preferences versus other considerations that they might have will, my guess would be, will vary in the future. And I think that some factors that could drive that would be your personal views on things like paternalism, how comfortable you are saying, "I know I have really good information and that I'm in a position to help someone make a decision," or whether you're saying, "I don't think I should have any role in deciding what someone else should receive, and that that is wholly up to them." And we see this with GiveWell's staff, and also with GiveWell's group of donors who use our research. I know that a lot of the donors that I've spoken with, that give to GiveDirectly of our list of nine Top Charities choose to do so because of paternalism and concerns about paternalism feeling particularly strong for them. And that that could be a reason why a GiveWell staff member or GiveWell supporter today decides that they would value cash over something else. You could imagine taking beneficiary preference responses and applying a similar logic to how you incorporate those and how you make those same trade-offs.

52:07 S?: Yes, so back to one of the other points about the setting between the nets and money. Do you guys have a consistent unit of related, or also what... How do you decide that this particular value is good, right, so in different parts of the world where the importance of giving different people different things aren't considered a good? So, is there a consistent unit of doing? Is it like good points or effective things?

52:42 CH: Yeah. So, do we use a standardized unit of good that we look at when we're coming out with a bottom line, and how do we think about the fact that what we think might be a unit of good would be different than what someone else would consider? And so in terms of the specific GiveWell approach, we use arbitrary values to make those trade-offs, so if you actually go into the cost-effectiveness model and look at the tab where individual staff members make these trade-offs and it's public on our website, so anyone who's felt inspired by this conversation to go in and see how we do it, or to put in your own values, that is on our website.

53:20 CH: So the way that we do it currently is that we're just comparing them relative to one another, so we're not saying we have a bottom line unit of good that we're translating all of our Top Charity outcomes into. But we're saying, "I value income twice as much as health," and that gets plugged into all the other assumptions that go into our cost-effectiveness model, such as how effective you think bed nets are, how much insecticide resistance is a factor, what you think the duration of benefits of cash transfers are. There are many different factors beyond these moral trade-offs in our model, those all come together into a bottom line outcome where we compare their cost-effectiveness relative to one another, and we make those comparisons via each individual outcome of a Top Charity that we recommend, so that is a very abstract sentence.

54:08 CH: So what that means is, if you go into our cost-effectiveness model, it will say, this charity costs X dollars as much, or, sorry, the relative benefit of this charity is X times as much as cash transfers or X times as much as malaria

nets. You could say the Against Malaria Foundation is 6 times as cost effective as GiveDirectly and that the cost per outcome is equivalent to... And then you can choose any of the outcomes on our list. And say like...

54:39 S?: It's really relative for the donors.

54:42 CH: Yeah.

54:43 S?: So it's like you get this amount of good points per dollar for this donor.

54:47 CH: It's per outcome, so you can say like this outcome... If this Top Charity were implementing a program with a different outcome, here is the cost of that outcome, so that you can look at it along all of those lines, and it might be easier to look at it because I think I'm not explaining it in the most clear way out loud. But in our cost-effectiveness model, you can see that this is the cost per outcome as good as averting the death of a child under five. Cost per outcome as good as giving someone \$1,000; and then you see those values for each of our Top Charities, and you can select the different outcomes that you're making those comparisons of, but we're not explicitly publishing like this is how many units of good you accomplish. We're putting it either into relative terms comparing our different Top Charities or putting it into the unit of a particular outcome that is translated via the cost-effectiveness model. Come see me after and I'll pull it up, because I think I'm not explaining it as clearly as I can out loud, but we can look at it on the sheet.

55:49 S?: [55:50] ____.

[overlapping conversation]

55:53 CH: I think we'll do one more question, because we're almost at an hour, but then we'll have lots of time to keep talking after that. So, yeah.

56:00 S?: To also piggyback on that question, I know in a lot of development studies, I always hear quality adjusted life year, the kind of metric used to compare different interventions across different... Like, whether it's health or whether it's cash transfer or whatever. Is there a reason why GiveWell wants to kind of look at an alternative to that and also to what extent is framed in the research?

56:22 CH: Yeah, so, the question is that quality adjusted life years and disability adjusted life years can be another kind of common metric that's used in global health, and why doesn't GiveWell use those? We actually used to use disability adjusted life years or DALYs in our cost-effectiveness model, but what we found is that GiveWell staff were ultimately using them in an unconventional way, where we were weighting the death at different ages differently to what the standard global health approach was for using DALYs, so it was no longer useful to us in making the trade-offs that we were making in the way that GiveWell staff were making them to use that metric, that common metric.

57:02 CH: And so I think it was just in 2017 that we switched away from doing that into the more, the sort of arbitrary, you could put in whatever baseline number you want and compare it according to your own values to get your own bottom line outcome, which is why we switched away from DALYs. The short answer is just that we were no longer using them in a conventional way, so that it wasn't helpful to us to use them, since it would be more confusing than illuminating to look at them, if you were familiar with DALYs.

57:35 CH: So I'm going to pause here, since it's been an hour, and say how much I appreciate all the questions and everyone here, and your interest in and support of GiveWell. It's really, I love coming here and seeing everyone and talking to you all. And I'm planning to be here for quite a bit longer, so I would love to continue the conversation one-on-one, and then I'm also going to ask Tracy to stand up, and Dan and Caitlin. Tracy, Dan and Caitlin all work at GiveWell, they can all talk a lot about the research that we're doing and the projects that they're working on, and would really love to meet you. So, strongly encourage you to chat with them as well, about, they can talk about any of these questions that I was just talking about and also many things that I don't know about, so you should definitely talk to them as well. We have this space until I think about 9:00, 9:30, so hang out, please. There's more pizza. And I hope that we can all meet and get to know each other. So thank you again for coming. Really, really appreciate it.