

GiveWell NYC Research Event June 20, 2016 – Top Charities

[background conversation]

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.

00:11 Speaker 1: Other thing is just that we know we produce a lot of written content. We don't expect everyone here to be up on all of it. Part of the reason we have these events is so people can engage with our work in a personal way, in a more informal way, so don't be afraid to ask questions even if the questions are basic. This is, to us, an alternate way of engaging with our work, it's not supposed to be kind of, you read everything, now what questions do you have? And that's why I'm gonna give a summary of everything we're doing. With that all said, the first thing that I'm gonna talk about is our top charities, GiveWell top charities. And I'm gonna start with the Against Malaria Foundation.

00:49 S1: GiveWell tries to find proven, cost-effective, scalable charities. We try to find the best we can, the best options from individual donor to do as much good as possible with each dollar that they give. And currently, as of December last year and still through now, our number one is the Against Malaria Foundation. They distribute insecticide-treated bed nets that people can use to kill mosquitoes, block mosquitoes, avoid getting malaria. These bed nets are associated with a large body of impressive evidence that demonstrates that they save lives and reduce malaria very cost-effectively. So it's like \$5 a net. Net lasts a couple of years, covers a couple of people, oftentimes, and overall, we estimate that it's in the \$2,000 to \$3,000, maybe \$3,500 range for life saved, and that estimate changes a lot because we try to adjust for all sorts of things, which I'll be getting into.

01:42 S1: In the past, we've recommended the Against Malaria Foundation a lot of the time, multiple years. There was one point at which a couple of years ago, we suspended our recommendation because we felt that they were not moving money out the door fast enough that we were excited to see people give there more. In other words, they have been sitting on a large pile of reserves for a year or two, and what the AMF has to do is they have to get the money and then go negotiate with African governments, with other funders, to put together a deal to have a distribution in a country to cover a high risk of malaria zone. And so, they need to make deals that meet their desires in terms of the cost per net, in terms of proof that the nets were delivered, which is something that we think they stand out on, in terms of follow-up data to see if people continue to use the bed nets as intended, if the nets stay in good shape and they need to get all those terms agreed to. And there was a time at which they kind of weren't finalizing those deals fast enough and they had been sitting on a bunch of money for a while.

02:46 S1: Recently, that has become a lot less of a concern for us. And a general theme of tonight and of GiveWell, is that we do try to discuss our top charities in just very frank, balanced, neutral terms. I will tell you what I think is good, what I think is bad, what we're happy about, what we're disappointed about. That's going to be a different tone I think from what you'll normally hear from people involved in the non-profit world. Our whole game here is sort of that we're not fundraisers, we're not marketers, we're just saying how it is, and when we estimate the cost per life saved, that number is higher than you might see from, a thing you get in the mail with a child's face on it, because it's not a fundraising page, it's our best estimate. And it's us trying to account for all the complications that we can come up with.

03:31 S1: So AMF, we have seen a lot of improvement on it's stability, put together deals and in particular, late last year when we recommended them, as a result of our recommendation, a lot of money went to AMF and a lot of it was from Good Ventures, which is a major donor that works closely with us. A lot of it was from individuals. It was about \$40 million total. This year, AMF has already signed two deals totaling about \$37 million for bed net distribution in countries that it has not previously worked in. It's in discussions about other distributions. It currently holds still about \$23 million on hand, but it's very likely to have to slow down its work and put the other fewer deals, if it doesn't get a lot more money pretty soon. So we think it probably is gonna have to slow down, and we would need another 11 million soon to move as quickly as it could, and we have no reason to think that's gonna happen. And it could use perhaps up to another \$18 million during this calendar year. So the picture with AMF has changed in some sense, in a positive direction in that sense that they are putting together more distributions.

04:41 S1: And in particular, one of the things that we used to be concerned about, that they prefer to pay only for the bed nets themselves and are reluctant to pay for distribution cost and monitoring cost, that has changed a bit. And they have been paying for some non-net cost in putting together those deals. We've been very happy to see that. On the flipside, we have been a little bit disappointed by monitoring reports that have come in from their work in a very difficult country, DRC. The reports came a little later than we hoped and the results were not... We can't talk about them fully yet, but I think overall, in terms of the quality of the data, in terms of the timeliness of the data, we've been, I would say moderately disappointed. It's not a game changer, it's not gonna change our recommendation of AMF, but

that is another factor. That is another update and that's important because it's our top charity.

05:35 S1: Another thing to know about AMF is that the global fund which is kind of the biggest international multilateral funder, is holding a conference in September where governments and other big funders are gonna be asked to make three-year pledges. And our early read is that there may be a lot more money for bed nets from other sources in 2018, than there is in 2017. So that has some interesting implications. That could mean that donations to AMF that are done soon are more valuable than donations that are done at the end of this calendar year.

06:10 S1: Most people give it the end of the calendar year but if you're someone who has some flexibility that might be something to consider, because there may be less of a global bed net gap, less of a global need in 2018 than in 2017 and AMF tends to work pretty far in advance. So that is something to know. It also does potentially affect cost effectiveness because you give out these bed nets and you hope they last two to three years. And so, if there's gonna be another distribution anyway in a couple years that could potentially reduce the bang for the buck by some relatively small factor.

06:42 S1: So that's a thing to know. And then we have done some other research on other things related to bed nets 'cause the stakes have gotten pretty high with AMF sending tens of millions of dollars there on our recommendation, so we want to kind of nail down everything we can. We did a series of little case studies on what happens when AMF is working on a distribution and they decide not to fund it. What happens? Does someone else just come in and fund it or do things get delayed? And the usual story was that things got delayed on average more than a year and so we kinda walked out of that case study situation, feeling that, not a hugely changed view and largely feeling kind of reassured that AMF really is distributing bed nets that probably weren't gonna get distributed otherwise.

07:24 S1: And then another thing we did is we have been very aggressively keeping up with the state of knowledge about insecticide resistance. So, over time, through a variety of mechanisms, mosquitoes may get better at avoiding the insecticide and it not being affected by the bed nets. And those changes could be genetic, they could be behavioral, but they could worsen the effectiveness of bed nets. Unfortunately, it is really hard to get data that is really useful about how bad resistance is and about how's it's affecting the effectiveness of bed nets. You can go to an area, collect mosquitoes, expose them to insecticide, and notice they're taking longer to die than they used to. But you don't really know how that's affecting the impact of distributing bed nets in real life.

08:07 S1: And so, there have been... There are a couple of high quality studies going on that are gonna be more informative than anything else that has happened before, but for now, our picture having caught up on the insecticide resistance literature and on the little case studies, is in the same range of where we were in terms of the cost per life saved and in terms of how good we feel about AMF. So bottom line, we are less concerned about AMF's ability to move money out the door, we've had a monitoring... Maybe a little bit less excited about the monitoring though that's always an ongoing thing. And we're probably not gonna make a big push for this but I think it is worth considering that there's an argument that if you wanna give to AMF, it's better to do it soon than to wait for December this time around.

08:51 S1: So that is AMF. Quick update on our other top charities, there hasn't been a ton that has changed. But SCI, the Schistosomiasis Control Initiative, they put together government programs that deworm children. Basically, give them pills that kill common parasites. And an issue we've always had with SCI is we've always felt it runs a good program. Deworming is a good program, it's very cheap, it's pretty simple, we think it generally works. It is a pretty good value for money. It's roughly comparable to bed nets though, in a very different way because it's kind of... For us with deworming, there's a good chance that deworming actually doesn't have a noticeable positive impact, and there's also a good chance that it has kind of an amazing impact per dollar.

09:35 S1: And so, it's accounting for those possibilities, though we think on balance it's about as good as bed nets. But SCI as an organization we've always had trouble communicating with them and we've never really felt that we understand them as well as we understand other charities, that we understand their financials in particular has been a real struggle to just understand where their money's going and where it's coming from, and what the situation is. They're at a university, and so they have these co-mingled financials that are not easy to make sense of, and that we have not had a great time trying to do that.

10:07 S1: This year we said we really needed to see improvement on that and we did see improvement. We've gotten basic financial info that we've had for other charities for awhile but we're really glad to have it for SCI. We also saw a couple of errors in the info that didn't make us feel perfect about the situation, but we do think SCI is improving on that front. We think it's still pretty far behind our other top charities in terms of how we feel about the organization. And it's

our clarity of understanding what they do and their ability to communicate effectively with us. But probably they're going to end up in a similar place on our top charities list, overall.

10:45 S1: And then, GiveDirectly gives out direct cash transfers to very low income households in Kenya and other countries and they are kind of at the opposite end of the spectrum where we do think, our best guess is that bed nets and deworming are a better value for money. They help more people for less money than just giving out cash. But we think GiveDirectly is really topping the list and by a substantial margin in terms of being an impressive organization that has shown an ability to scale, that has been creative, that has really come a long way in a short time.

11:19 S1: One of the cool things about GiveDirectly is I believe all or nearly all the money that they're dispersing in cash transfers is also part of experiments so they're always trying two different approaches, randomizing them, comparing them, always trying to generate knowledge as well as generating positive things by giving money to very low income people. This year an example of a couple of new projects they have, they've working on a 25 million dollar initiative to test a basic income guarantee. So this would be instead of giving people a lump sum of... A one time lump sum which is intended to encourage investment, they would give people a regular income and test it and run a big study on it and see how that... What they can learn about how people respond to that.

12:03 S1: And that's obviously a hot topic in a lot of different parts of the world, and it is a different topic in America than in the developing world, but it still, I think, will be very interesting. And then, they're also working on a partnership with a really major funder to do a series of studies comparing normal, sort of, standard aid interventions to cash transfers. So, in other words, if we're doing an agricultural training program, where we try to help farmers get better yields from their land, is that better than just giving them the money that it would have cost to run this program? I think that's probably a question people should have been asking for a long time, but now there's going to be a series of studies asking that, very directly, in a randomized setting. And we're pretty excited about that.

12:46 S1: GiveDirectly did run into a snag in one of the areas they're working, Homa Bay, Kenya. It looks like local leaders may have been advising people not to accept the transfers. They ended up moving to another part of the county. And that's the kind of thing that happens in these situations. They're working in a lot of different places, and they've really scaled up quite impressively. But the bottom line is, we feel similarly about GiveDirectly, they're impressive, they have extensive room for more funding, and that is where we're at. And then, the final update is on Deworm the World, which is another organization that does deworming, treating children for parasitic infections. And they are working on expanding to other countries, but they are not funding constrained right now. So, that is not a major target for donations at this time.

13:33 S1: So, that's kind of the top charity status. Overall, we've still got Against Malaria Foundation in the number one spot, all things considered. People might wanna consider moving their donations up. That said, we are not recommending to Good Ventures, the major donor we work with, we're not recommending that they close the gap. And the reason we're not doing that is because Good Ventures made a major grant to AMF last year, and that grant was intended to, kind of, be a large commitment, but also be the end of the commitment for that year. Because we don't wanna get into a situation where Good Ventures always fills 100% of every funding need they see, and therefore, that creates a situation where other people have no incentive to give and do their part.

14:15 S1: So, that's a longer topic, we had a very extensive discussion about it, I think, probably at the last event, and also there's been extensive discussion on it on our blog. But that said, I'm happy to answer questions about it. I've just been talking for awhile. So, that is the top charities update. And now I'm going to take questions for a bit before talking about some of the things that we're doing to find more top charities in the future.

14:36 Speaker 2: One thing I just wondered. For, like, SCI and AMF. If those two organizations didn't exist, are there similar organizations that are a close second in the same field? Like, in the bed net field, and in the...

14:50 S1: Sure. I'm going to be repeating questions, that's for the recording. So, if SCI and AMF didn't exist, would there be a natural next place to turn for bed nets and deworming? Well, in the case of SCI, there's Deworm the World. So, there's two recommended charities that both do deworming, one of them is not funding constrained. And we don't have anywhere to go after SCI. I mean, we're actively working on it. We're looking at a couple of contenders. But I would say there's no obvious place where you can write a check and expect more people to get dewormed beyond those two. AMF, I would also say no. I would say that there are other organizations that will, you know, that have a mission of delivering bed nets.

15:31 S1: But if you want any kind of reporting, confirmation that the bed nets got delivered, check ups on whether people are using the bed nets, checking up on Malaria case rates, where the bed nets were distributed, we've gotten just zero engagement from anyone else on the idea of collecting and sharing that data. So, you could... There's a group called Nothing But Nets, they work with the UN Foundation. By mission, they do bed net distribution. You could write them a check. You could hope for the best. But that is all I can say about Nothing But Nets, because they've declined to participate in our process.

16:05 S2: Other people can go if...

16:09 S1: Sure, you should go ahead, for clarification. Yeah.

16:16 S2: What kind of reporting does AMF do on bed nets, and along those lines, is there reason to think that other organizations that deliver bed nets have... That reporting would come back worse than AMF does?

16:35 S1: Sure. So, what reporting does AMF do, and is there reason to think that it's important? I would say... So, AMF does a variety of things. They look for proof that the bed nets were delivered, often with pictures. They look for post distribution... They look for pre-distribution surveys of a certain thoroughness, checking who needs a net, which I think AMF does that more thoroughly than a lot of other distributions and that can reduce wastage. And they also look for post-distribution follow ups, to make sure people got their nets, and come in later and see if people are still using them, and if they're still in good condition. It's very hard to compare this to other bed net organizations, because if you don't measure it, we don't know what the situation is. So, I think it's somewhat of a matter of judgment. If you'd rather write a check to Nothing But Nets, and hope for the best, there's an argument for doing that. But there's not a tremendous difference here.

17:27 S1: That said, with AMF being as funding constrained as it is, there was a time when we were kind of saying, "AMF is having trouble getting money out the door. Our recommendations are moving a lot of money. Let's look for the next bed net organization, even if they're not as strong on the monitoring." But we've dropped that for now, because I don't see any reason to do that at this point. AMF has plenty of need for more money and so, I think it kind of dominates the other options. But again, it is a matter of a judgment call. I mean, those distribution and monitoring costs are a small fraction of the cost, so it's not... I don't think you would get a big boost here, cost effectiveness, by cutting back on 'em, you would just get even less information than we have about how it's going.

18:07 Speaker 3: Do you know offhand roughly what percentage of their process they are?

18:12 S1: I think it's under 10%.

18:14 S3: Perfect.

18:15 S1: Yep.

18:17 Speaker 4: This is, I guess maybe more of a question for like Open Philanthropy, but what is the cost benefit versus say, funding initiatives, to say, for vaccine or [18:27] _____.

18:34 S1: Sure. So what is the cost benefit of initiatives like funding development of new vaccines versus the stuff we do, which is very straightforward and linear? A general theme of GiveWell, is that GiveWell is a product in some sense for people who want to give, people who wanna do as much good as possible, but also who don't kind of have a lot of time to decide and would like to give based on verifiable information that is straightforward to explain and straightforward to demonstrate that the claims we're making are reasonable.

19:06 S1: And so GiveWell very explicitly and very deliberately has a bias toward things that are linear in this way. Things that you can look at where your dollar goes, what it's paying for, you can prove it, you can estimate how much good you're doing per dollar. And I have to say, all this stuff, I mean if anyone has a picture that our top charities are kind of iron-clad cases, they're not at all. There's a lot, even when we go... Have this massive bias toward measurability, there's still a lot of problems and a lot of things we don't know and a huge amount of guesswork we have to do, and when we estimate those costs per lives saved, we publish our spreadsheet with all the different numbers you have to guess at to get to that number, and anyone who's ever downloaded that thing knows there's a lot of uncertainty here, and those numbers can come out a lot of different ways.

19:50 S1: So I don't wanna overstate the case, but we do look for the charities that look like they have the best case for this kind of linear demonstrable verifiable impact. What happens when you drop that requirement and you just say, "I wanna do the most good possible." According to me and according to my best guesses, even if there's much more debatable and unknowable things here, and I think that is the mission of the Open Philanthropy Project, like you said, which is the next session. I will say that a couple things happen that I think are worth knowing. One thing that happens is that the further you go down that road, the more your opinions and your views, I think I predict, are going to diverge from the rest of the world's. So, when I talk about Open Phil... When I talk about the top charities, I think a lot of people are listening and saying, "Yeah, that sounds good. Those sound like they're helping people a lot, but is that really the best thing you can do?"

20:43 S1: But when I talk about Open Phil, I think a lot of people are gonna hear things and think, "I guess that could be the best thing you could do, but that sounds totally crazy to me, and actually it sounds horrible." And so that I think is it, that's the spectrum, and we're kind of more interested in the two ends of it than in the middle. We're kind of more interested in trying to do the most verifiable things that we can pitch to anyone that anyone can believe in or just trying to do the most good possible, even if it looks like really wacky to people who are not up on the issues, than in that middle ground so far, and maybe later we'll decide there's a case for the middle ground. But I think that is one thing that happens as you go in that direction, and so if we were to look at vaccine research, we'd all be just massively guessing at the odds of getting a new vaccine, and about what the scientific prospects are, and whether this makes sense, and the more vaccine looks like it's easy to make and definitely can be done, the more it's already gonna be funded. And the more things we find that are like, "Oh, we can make a vaccine for this, maybe." It's gonna be a big "maybe" and it's gonna be a huge judgment call. That's one thing that happens.

21:47 S1: Another thing that happens is that so far, my take is that the top charities look pretty good compared to anything, and so having spent several years, thinking about how to be the most ambitious possible and look all over the place, we've looked at all kinds of things from political advocacy, to scientific research, to global catastrophic risk prevention. Are there other options that I can make a case for better? Absolutely. Are there other options that according to certain world views and certain biases and certain philosophical presuppositions are a zillion times better? Yes there are. But are there interventions that I think are clearly better than our top charities? Not so far, and we've spent a while looking for them, so I think that is another thing. And so when I give personally I give to top charities, and I think it would be even more true... I think with me, it's kind of a strange case because the Open Phil project, we kind of generally tend to fund what we recommend, and so I'm partly into the top charities 'cause they are the thing that isn't already funded. But it's also, if I had a different job and I didn't think about this stuff as much, I would especially wanna give there because I think it is an easier made case and I don't think you miss out on a clear, obvious massive win by sticking with this measurable stuff.

23:04 S1: I think when you're working on the developing world and you're helping people with such low incomes who have such basic needs, there really are, in my opinion, kind of very unfortunate, and it says something negative about the world, but there really are some really basic needs that people have and some ways in which money can help very straightforwardly that I think are very good and not really easy to beat by looking elsewhere. Yep?

23:35 Speaker 5: How many NGOs do you evaluate each year to join the list of most recommended?

23:42 S1: Yep. How many NGOs do we evaluate each year to join the list of most recommended? So this has changed a lot over time. When we started GiveWell, we invited I think a few hundred charities to apply for grants and it was kind of every reasonably large charity we could find that worked in relevant ways. And we kind of narrowed the field from there and in 2009, and I think in 2011, we did these big sweeps where we looked at a lot of charities' websites and we had an early filter, that we were kind of looking for anyone who was publishing either suggested things on their website that showed they might be evidence backed, or someone who's working on intervention with a good independent literature behind it.

24:23 S1: But having done that several times, we do feel we've hit diminishing returns and I'll get to this a bit when I get into the mid session talk on future top charities, but we sort of feel like at this point, we've defined the criteria for top charities, we have learned that usually, when there's really strong evidence behind an intervention, that evidence is academic literature, is out in public, it is known, it is findable, without going to the charity. I think when we started, we imagined that a lot of charities have great evidence buried in a file drawer somewhere, that they would give us if we asked them, we don't really believe that anymore.

25:00 S1: And so, now the process looks different, and now the process is more about knowing there's a certain number

of interventions that have the evidence and keeping a really strong eye on the charities that focus on those interventions, and checking in with them to see where they are and whether they'd be interested in our process. And so, we often look at three or four charities a year and we look at them really deeply and we really try to dig and understand everything that's going on. And that feels like a more realistic way to find charities with these qualities. So most charities, I can look at in 10 minutes and say, "This is an interesting program but it is not one that is supported by academic literature in the same way that our top charities are, or it is not one that is going to be as cheap as bed nets and you're not gonna get as good value for money in terms of helping people." And so that is the way our process is now. We've got the criteria pretty squarely defined, and that means the pool of plausible contenders is limited. Yup?

26:01 Speaker 6: I think people are often confused by too many choices in a lot of different fields, [26:07] ____ on this, but are you concerned that maybe by limiting the number of choices, maybe not having a diversified option to invest, that you're maybe kind of refusing some [26:19] _____?

26:21 S1: Sure. Are we concerned that we're offering too few choices to people and therefore discouraging giving? I think it's definitely a concern one could have about our work or it's more... I think I might describe it less as a concern about our work and more as an opportunity for someone else. I think this is kind of a marketing decision, and kind of an audience and niche decision where our view is that the people we're trying to help are the... The people we're ultimately trying to help are people who are receiving bed nets and deworming medication, but the kind of customers of our product or the clients are people who are saying, "I wanna give some money and I wanna do as much good as possible." And that is just a very different mentality to be in than a lot of other donors who might say, "I already know what charity I wanna give to, can you tell me if it's legit?" That's not our mission, that is there's other groups that do that. There's other people who say, "I wanna give to this cause. I want it to be micro-lending, I want it to be in Ecuador. Do you have a charity for me?"

27:23 S1: That is not... Also not our audience and also not our product and there I've seen some attempts to kind of do that. So I would say it's more just a decision about where to focus. We focus on an audience that is similar and relevant respects to us and what we wanted when we built this product. And so as a member of that audience or a hypothetical member, I would just rather someone really nail what the best opportunities are and I'd like to give there. And I think different products could serve different people and I think it's great that people have efforts to do them but it's just not the strategy we've chosen.

28:00 Speaker 7: I read every once in a while about area and mosquito control strategies that people are coming up with. Like there's the thing in Brazil where they irradiated a bunch of male mosquitoes or whatever. Do you guys evaluate that? Is there a sense that that's coming up in cost effectiveness against bed nets as a strategy for malaria reduction or is it still a long way off?

28:19 S1: Sure. How do we feel about area control, different ways of controlling mosquito populations and are those competitive with bed nets? Will they be? There's a lot of different things under this heading, so, indoor residual spraying has been around for a long time and that's where you kind of, you spray insecticide inside houses and you hope that kills a bunch of mosquitoes. That's something that I just believe is a lot smaller part of the malaria control picture than bed nets, and we would... We don't know charities that focus on it and we don't think we'd be able to get the same volume of donations moved to it. It is something we've thought about from time to time but there's no charity doing it, we don't think it's better than bed nets. Though it often is a good part of the malaria control strategy. There's some strategies that have been used I think very effectively in different settings. So I think in America, getting rid of malaria here whereas it used to be a big problem and is now not in America at all, there used to be a lot of swamp draining.

29:15 S1: And so there's a lot of things, we've got swamp draining and fish that eat mosquitoes, I've heard something about that and there's a lot of methods that have historically been used I think more in areas that we're trying to get from low levels of malaria to zero rather than these very endemic areas. So my... It's not something we've ever done a formal report on. We did have a lengthy conversation with a malaria control expert that we never got to publish the notes from. But our impression is that bed nets are the kind of... They are the thing you can do without infrastructure at great scale across a large distance, get good effectiveness, and we haven't seen any real competitor to them in terms of a charity.

29:54 S1: Now what you may be talking about is more future control efforts and that could include things like there's been a lot of headlines about genetically modified mosquitoes that might be malaria resistant, or that might just be reproductively non-viable. That is the kind of thing that could conceivably be a fit for Open Phil, but we're not gonna say, "Hey, our top charity this year is taking a bet on the possibility that there can be genetically modified mosquitoes 10 years from now, and you should support them." That's where we've decided not to be as GiveWell. Yep?

30:25 Speaker 8: So, as a donor, if your top recommended charity is AMF, and... I really appreciate the important work you do, but I tend to, given in the past months to the AMF, but I tend to like GiveDirectly more lately, 'cause I just feel like in the case of AMF, people who are being saved, this sounds horrible, but they were going to die. And I feel like that's not as bad, as someone who's trying to do good in the world, I feel like that's not as bad an outcome as somebody living in poverty. So, I wonder how you weigh that, and basically, I guess I'm saying, how would you convince someone like me to donate to AMF or to GiveDirectly?

31:15 S1: Sure. How should we compare AMF to GiveDirectly, when AMF is largely preventing deaths, and GiveDirectly may be improving quality of life? So, a few comments along these lines, one, I think this is a tough question, and it's not one that I would expect everyone to agree on. And so everyone knows, the way we do our cost-effectiveness estimates is we set up a spreadsheet, the spreadsheet has a bunch of places to put in numbers that can be debated. One of the numbers that can be debated, and actually this year's edition of the spreadsheet will highlight this more, is how you wanna value averting a death, especially averting an infant death. Because I think people have very different intuitions. Is it as bad for someone to die at age 18 as it is at age three, or age three months? People have different intuitions about those things. And so you can put, 'How do you want to value averting a death' versus 'improving a year of someone's life'. And there are multipliers that, really, it's just using intuition. It's like would you rather double someone's income for a year, or give them an extra year of life?

32:14 S1: Then what we do is, all the staff members put in all their best guesses, and we have a lot of debates and discussions about 'em, and then out of that big spreadsheet comes a number. So, when I put in all my numbers, I get result that AMF is five times as good as GiveDirectly or whatever it is, and then we take all the staff inputs, and these days we'll have 10 of them or something, and we'll list them all, and then we'll take the median. So, that's our official cost-effectiveness estimate. If you have different values, there's different ways to handle this. One, you could download the spreadsheet yourself and put 'em in there. You could copy another staff member's inputs, because some of the inputs are a little bit esoteric. Though we try to explain 'em in the spreadsheet, I think we could probably be doing it better. Another thing you can do, is you can just go with your gut. And you can say, "Hey, I believe I'm probably not gonna value averting deaths as much as improving quality of life, so I want a quality of life charity."

33:06 S1: The final thing I'll say is that GiveDirectly is a quality of life charity in that sense, but so are SCI and Deworm the World. Deworming is also... The benefits of deworming are just basically not at all coming from saving lives. The parasitic infections are fatal so rarely, and they're not the kind of thing that are affecting our calculations, so that's a quality of life thing too, that's basically an income boost, just like GiveDirectly. Is the benefit the world aware of? And so it's kind of, how do you trade off the uncertain future income boost versus the certain short term, maybe, smaller income boost?

33:41 S1: And then, AMF is good to keep in mind. AMF is not just a death averter. AMF reduces cases of malaria. And malaria can be debilitating, it can cause permanent damage, it can stress out the health system, it can have developmental effects, just like parasitic infections can. It can weaken people, it is associated with anemia in pregnant mothers and infants. So, you get a pretty solid quality of life boost with AMF, but if you are counting the deaths averted as zero, you'd probably like deworming better, because it's cheaper for the same kind of boost. So, that's the answer there.

34:18 S1: I'm going to go into the second phase of talking about future top charities. And then I'm gonna take more questions, but I'm happy to take questions at any topic, so it won't just be limited to that. But, one of the big questions that has come up over the years for GiveWell, is, we have the same top charities every year, or some approximation of that. And I think, on one hand this isn't necessarily a problem. Our goal is not to be exciting and fun, and have a new top charity every year just for the fun of it. But, on the other hand, it does make you wonder, "Can we find more good giving options, and what can we do?" And one of the things we've learned over the last few years is that the current method is probably not gonna work.

35:03 S1: We feel like we probably are aware of all the serious contenders to meet our criteria, and we've gone back to many of them, several times. And this year, we tried to start doing another round of this, of going looking for the contenders, and going to them and asking if they want to participate. And it was like we were looking at the same organizations, getting the same lack of interest in being a recommended charity, which surprised us. Because the advantages of being a recommended charity are much higher than they used to be.

35:30 S1: At this point, if you're a GiveWell-recommended charity, you are probably going to raise all the money you

can, all the money you can spend, unless that money is getting above \$5 million or \$10 million a year. And even then you might raise it all. AMF got something like \$40 million, and those numbers are going up. So it seems like the incentives are there to participate in our process, and so my best guess is that charities that don't participate in our process, probably, in many cases, do not have the data we were looking for, cannot demonstrate the things that we're asking for or if they can, they probably very strongly don't need more money or something.

36:11 S1: And so that approach just looks less promising to us than it did. And the question is, what else can we do? And I think one interesting fact is that our top charities, none of them are super old. So I think SCI is the one that's been around the longest or maybe AMF, but AMF was very small compared to what it is now until we started recommending it. And so, AMF was like under a \$1 million a year organization at the time we recommended it. Again, we attributed \$40 million of revenue to it last year from our recommendations. SCI has been around since maybe a year before we started, GiveDirectly was founded several years after GiveWell started and Give It To The World also. And so these are relatively new organizations, but the question is, could there be future top charities and can we make them more likely to come into existence by making recommendations and grants today?

37:07 S1: And so the work we do on this we refer to is GiveWell experimental, and it's a little bit of a hybrid or a middle ground between GiveWell and Open Philanthropy. The work is done by GiveWell staff and I don't really work on GiveWell. I basically work on Open Philanthropy, and so this is work that I'm a little bit less intimately familiar with, but it's GiveWell staff trying to figure out which charities and which interventions and which pathways could lead to a future charity that meets all the GiveWell criteria. But methodologically and in terms of the fund that we work with, it's more like Open Phil in that we're willing to take long-term bets, high-risk bets based on questions like, do we believe in the leadership here? Which is also a factor for top charities. And then we generally make recommendations to Good Ventures which is a major donor that we work with, and Good Ventures has been funding things that we think can turn into future top charities. And these come in different shapes and sizes so I'll give a few examples of this.

38:04 S1: One grant recommendation we've made recently is to a group called R4D that is trying to improve access and execution of pneumonia treatment. They believe that pneumonia is neglected relative to the "big three diseases" which is HIV-AIDS, Malaria and Tuberculosis. They're starting in Tanzania, they're gonna try and increase coverage of antibiotics, improve health worker practices, improve distribution, get more accurate diagnosis. This is obviously a more complicated thing with more moving parts than distributing bed nets and it could be very hard to evaluate which is a worry of ours, but we do think there is at least a possibility that this will work, that we'll be able to see it work, that they'll improve pneumonia access, save lives for a value for money somewhere in the range of AMF and that would be great.

38:55 S1: And then it may scale up to the point where we're able to see the past and say, "This worked," and we can get more of it and they are a top charity. So that's one example would be a group that is trying to scale up and trying to prove themselves, and could directly become a top charity. And similarly, we recommended a grant to New Incentives which does conditional cash transfers which is a program that we think is very promising, we just don't know any charities operating at scale that do it. And so, we're funding this kind of early stage charity that gives out cash transfers in exchange for people proving things like their child going to school or getting certain vaccinations and you get a cost effectiveness boost there. So those are some things we're doing. We're also interested in supporting work that leads not just to new top charities, but to new priority interventions. So one of the big filters, one of the big ways we narrow things down is we only want to look at charities that are working on proven programs. So if we can fund more programs to become proven, then that helps, too.

39:58 S1: And so along those lines, there's this group, IDinsight, they do decision-focused rigorous evaluations for decision makers. So it's more... There are a lot of groups doing randomized controlled trials, doing studies of aid programs, but a lot of them are academic and a lot of them have academic incentives, they try to work on interesting questions in an academic way. And that is different from looking for, really focusing on helping non-profits and policy makers and funders make decisions. And so IDinsight is a unique group that works on evaluations that are often not intellectually interesting in the same way as academia, but really valuable to decision-makers, and we recommended general support to them. So not even a particular project, just expanding their operations making them able to take more projects and maybe produce more proof of things that worked.

40:52 S1: And then the final thing we've done is there's this No Lean Season program which is done by Evidence Action, the same group that does deworming. And it's this interesting program where basically you provide financial incentives for people to migrate from rural to urban areas during a tough season. And there's evidence from randomized control trial that when you do this program, people will still migrate a year later even without the financial incentives.

And so that's an income booster and an information booster and that's something where there has been a randomized control trial and we're funding more of a pilot scale-up and hoping to get them to the point where they can be a top charity.

41:32 S1: So that's what things look like and we feel like this stuff is it's a little hard to say when it will bear fruit or if it will ever bear fruit, whether these things will ever turn into top charities, but it's probably our best guess path at how to get more. So with that for another five minutes or so, I'll take questions on anything about GiveWell and then we'll take a break and talk about Open Phil. Yep?

41:56 Speaker 9: Going back to the question earlier, so gene drives are new kind of experiment, it's new experimental technology that's only been demonstrated in laboratories. But there are these Oxitec transgenic mosquitoes that have been demonstrated in the field, in a few countries. I was wondering if that is something that you have looked at or managed to look at either GiveWell or Open Philanthropy?

42:26 S1: Has GiveWell looked at Oxitec mosquitoes which is a more mature form of genetically modified mosquito? Oxitec mosquitoes are not a fit for GiveWell, similar reason to gene drives. It's not just about the tech not being there yet, it's also about all of the policy and political questions about releasing genetically modified mosquitoes into the environment. I think that is often a big question about Oxitec. I'm not super familiar with the evidence, but I don't think it would be yet at the same level as what our top charities are. I think this general area is a potential area of interest for Open Phil, but not something that we've done a lot with yet. But it is something we'd pay attention to, for sure.

43:16 Speaker 10: Just going to go on the topic of new top charities, but more generally, could you summarize your impression of how the cost-effectiveness has changed in general over time and any [43:26] ____ adversity you might have over that story?

43:29 S1: Sure, how has cost-effectiveness changed in general over time? Things like our estimate of how much it costs to save a life, how has that changed? I think in our early years, we were way too optimistic in this... It's weird to call it optimistic. In some sense, the cheaper it is to save a life, the worse state the world is in, in some way. In our early years we had lower estimates, we thought it was cheaper to do good. At some point we kind of got deep enough into the issues that we discovered a lot of things we hadn't been thinking about, and a lot of concerns that hadn't been on our radar, and since then I feel like for the last few years, our cost-effectiveness has kind of gone up and down but not in a super-systematic direction.

44:15 S1: I expect that 10 years from now, I expect that cost-effectiveness of top charities will be worse because there are... The world's getting richer, the amount of aid is going up, the bed-net gap may end up getting closed at some point, completely. Deworming is a pretty cheap intervention and I won't be surprised if governments are funding the whole world to get dewormed 10 years from now, and it's a newer intervention so it hasn't been in the core government wheel house, but it could be.

44:45 S1: I would expect... There's other things, there's new technologies that can help control both parasites and malaria, and Africa is growing economically, and that helps, too. I expect some of these top opportunities to not be in existence anymore in 10 years. And I expect those cost-effectiveness estimates to be more in the range of GiveDirectly. GiveDirectly is a pretty good lower bound in terms of... It will be true for a very long time, that you could spend a very large amount of money if you want, giving money directly to very low-income people, and if you're systematically going after the lowest income people in the world, then even as the world gets richer, that opportunity is still gonna look pretty much the same, 'cause there's a big pool of people you could help, and you keep targeting those people and there's a lot of them. That's where I think those are going. Yeah?

45:41 Speaker 11: Are you able to break down your cost per life saved, say for AMF into a few key inputs, and maybe contrast that with how another organization might... Might price out cost per life saved?

45:55 S1: Sure, contrast our cost-effectiveness calculation with maybe a more typical thing you might have heard. I think the way I would do this is... Bed nets cost about \$5 and bed nets have some sort of relationship to preventing malaria and saving lives, and so a number... Well, actually, I think they're a lot less than \$5 now, but they kind of were last time as, I was engaging with this particular debate, I think they may be more like \$3, let's say \$5. A thing you will hear sometimes is "Send a net, save a life." That's the slogan of one of these charities. You'll hear this kind of tenuous, like, "Well you can save a life for \$5" which, sort of, implicitly means, for \$5 you can buy a bed net and bed nets do save lives. Our calculation is a little bit more... Well, is very different from that.

[laughter]

46:43 S1: And so the estimate is that... You start with the Cochrane review of the 20-something randomized controlled trials of bed nets of which I think five were looking directly at mortality and the pooled analysis of just kinda putting together a lot of different studies and combining them into one kind of virtual study has an effect... And I may get this wrong, but it's something like, one infant death averted for every 250 infants covered with a net for a year. So it's like person years of protection would be the metric. Obviously, that's where a lot of the difference comes from, because it's not like everyone who has a bed net would have died if they didn't have the bed net that year, that's more like a one in 250 thing.

47:29 S1: Then the bed nets are not all going to infants; it's universal coverage, which is how these things are always done, because you kill more mosquitoes that way, and because it's easier to get people to use the nets that way, and because everyone suffers from malaria, it's just the deaths tend to be in under-fives. Though we're not sure about that either. You get, let's say, 20% of the population, I think, is in that under-five range. And then... See, you've got one in 250, then maybe it's more like one in 1250 person years of protection of all ages is a life saved.

48:02 S1: But then you get average of about two years of covering two people each with a net, so it's like four years of protection for the \$5 that the net cost. So it's saves \$1 per person year of protection and we just got to \$1200. Then we have to put in some adjustments for insecticide resistance, we have got to adjust the threat of insecticide resistance for the threat of nets getting wasted and for a bunch of other factors that I'm forgetting but at least got within a factor of two or three, so that's the difference between those estimates. Yup?

48:37 Speaker 12: So it's Ryan, so I have two short questions, the first one is the welfare your GiveWell experimental is like, is it like taking off, like it's gonna be really big in five years like Open Phil, or is it like you are just dabbling in it? That's question one. Question two, I was wondering if you could provide a bit of clarification about why you think the charity that you mentioned where you give people incentives to move to the city when they're having off season is like likely to be a good one, because I didn't really follow what you were saying about people move anyway later, or something like that.

49:16 S1: Sure, okay. So number one, is GiveWell experimental going to become a huge thing or is it dabbling? I think that's not, we are not sure yet. But so far, it is dabbling, we don't have some kind of business plan that says this is gonna become a 40 person massive thing with a huge budget. We're looking for opportunities as we find them and if we find ourselves overwhelmed with good opportunities, we may have to staff up to handle more of them but so far, it's very incremental. Why do we think No Lean Season is a potential good program? Well, so the study I'm getting people who moved to the city didn't just show they came back the next year, it also showed actual positive impacts on their incomes, so by moving, people make more money.

49:56 S1: A lot of these charities really, the case is about reducing poverty and improving wealth, so it's that. I think the idea that they came back the year later is just like a positive sign that this program was adding some sort of value beyond just handing people money. It was giving people information in some sense that changed their behavior in a more stable way. So I'm now gonna call a break, we're gonna take about 10 minutes, and then I will talk about the Open Philanthropy Project. So thanks, everyone.