

## GiveWell NYC Research Event December 5, 2016 – 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.*

00:01 Holden Karnofsky: Alright I'm going to pull my notes and start talking about GiveWell. So the first session is GiveWell and GiveWell's top charities. So just as a quick refresher, GiveWell is an organization that tries to find the charities you can give to to do as much good as possible with your dollar. Specifically, we focus on organizations that are doing proven cost-effective and scalable work, so generally look for really strong evidence of effectiveness. We generally work on organizations fighting poverty, helping the worst off. We also look for a good bang for the buck, and that's actually why we tend to focus on the worst off. We look for ways to do a lot of good with a little money, and we look for organizations that are very kind of scalable, in the sense that if you write them a check, they will do more stuff. So someone giving out bed nets, we make sure that if they get more money, they'll give out more bed nets.

00:52 HK: These criteria are really optimized for individual and, in some sense, casual donors in the sense that if... GiveWell was originally started when Elie and I, Elie, my co-founder and I, wanted to give to charity and we were in the position of saying, "We have full-time jobs, we want to give to charity, we only have a few hours to decide where, we wish there was a helpful guide." And if that's the position you're in, I think GiveWell gives you something where we've really gone over the evidence, we've looked at things from every angle, we've published all the details to our website, and so you can give confidently to GiveWell's top charities, and that is really the goal of it. The thing I'm going to talk about next session takes a contrasting approach; so it's the Open Philanthropy Project and it's more focused on serving people who might be full-time donors, in a sense, and it is much more interested in kind of very high risk and very long-term giving, but that's not what GiveWell does. GiveWell is about things that we can make a confident and transparent and public case for.

01:52 HK: So that's kind of the difference in philosophies between the two, and so the main thing I'm going to do in this first little session is talk about our top charities. For several years, we've had the same list of top charities, there's four of them: Against Malaria Foundation, SCI, GiveDirectly, and Deworm the World, and then I'll get to more later. But for several years we've been looking for more top charities to add to our list, at the same time as updating and making sure we're staying caught up on everyone, and that no-one stays on the list who shouldn't. And this year, for the first time in several years, we did add three new top charities, so we went from four to seven. And now we have quite a larger number and so instead of going by charity, I'm going to try and organize this going a little bit by intervention, because a lot of our top charities kind of do the same thing. And then I'll talk about the big picture, just like how this year is different from past years, things like that.

02:49 HK: But I'm first just going to give a summary of what our top charities are, what they do, what the news is on them, and then I'll give some big picture thoughts and take questions. So, starting off with cash transfers. GiveDirectly is an organization that their basic mission is for every dollar you give them, they are going to get about 82 cents to people with extremely low incomes by global standards. So this is often a matter of going to the poorest parts of Kenya that they can find, and furthermore kind of looking at the people who do not have a metal roof, who have, let's say, a mud thatch roof. And so they're really trying to zero in on the lowest income possible people, and just give them cash, using this mobile money system where people kind of store their credits on a SIM card.

03:39 HK: And the basic thinking is there's a couple of reasons this is attractive. One is you're making your money go as far as possible by getting it to people that it will mean a lot more to than it does to you or anyone else. And another key philosophy behind cash transfers is that it's about empowering the people you're trying to help. It's about instead of us saying, "This is what you need," we're saying, "We have more wealth, we're going to ameliorate that to some degree, transfer the wealth, and then people can decide what they want to do with the wealth." So it's kind of all about empowering and leaving the decisions up to them. This is an intervention that I think just instinctively makes a lot of people really uncomfortable. Personally, if I had never heard of charity, I would probably guess this would be the first thing someone would have tried as a charity, but in fact, it's kind of the last thing.

[laughter]

04:32 HK: It's kind of a new weird thing that has become more of a topic in giving circles and development circles recently and it's a new organization, only a few years old, it's younger than GiveWell, and I think that's pretty interesting. I think a lot of people have a couple intuitions about why it's not a good idea to do this. I think a lot of people are afraid the money is being spent on, let's say, alcohol and entertainment. A lot of people just think, "Is that really the best you can do? Just give cash to people?" But cash transfers actually, partly, I think, because of the

skepticism, have been one of the most extensively studied programs out there in the development world, and I think the results are actually pretty encouraging. People have looked very hard for negative effects, for signs that a lot of the money is going to alcohol or things like that, cigarettes. Haven't really found it, have generally found that people spend the money in a variety of ways, a lot of times they might invest it. Buying livestock is an example of something someone might do if that turns out to be a good thing for them to do that can help them earn more in the future, or fixing the roof so that they don't have to... Or getting a metal roof so they don't have to constantly spend money fixing the roof. Or just buying more food and having better nutrition.

05:42 HK: And so the studies that exist, I think, have been pretty encouraging about the effects on people's lives, about the investment returns people get, and we think this is a good intervention. We do think you can probably do better, according to our best estimates. The other charities I'm going to talk about, when we try to calculate it all out, it looks like they're doing significantly more good per dollar spent. But, if we're wrong, if our calculations are bad, I think GiveDirectly is a pretty good deal. And I think especially for people who are especially skeptical, skeptical of studies, skeptical of calculations, skeptical of evidence, this is a nice simple intervention. I think it makes a lot of sense and I think your money can do really an enormous amount of good because the people are at very low incomes. So \$1,000 might take an entire family and be a year's income for them. And they get it all at once, which gives them opportunities to invest, and to pull themselves out of situations that they rarely get. It actually can be quite transformative and it is quite a great return on your charitable giving, in my opinion.

06:49 HK: The update on GiveDirectly as an organization: They remain the strongest organization, I would say of our top charities, just in terms of they have really excellent monitoring, we really understand what's going on, they're great at collecting data, they're self-critical, they're transparent, they studied themselves, and they're scaling really successfully. So in the next year they could easily spend tens of millions of dollars this way. There's very little doubt that they can absorb whatever money people choose to throw at them. So I would say this is the top charity with the least room for doubt, that you are doing a lot of good if you give there. Next up is kind of... Sorry, one more comment on GiveDirectly. They have been ambitious in a certain way, which is that they've often combined what they're doing with other things.

07:36 HK: Almost all or I think all of the cash transfers they give are done as part of experiments, where they might try two different ways of giving cash transfers and see which one is better. They're working on a partnership where they're going to be testing a more traditional charitable intervention against cash than kind of promoting this idea of cash as a benchmark. If you're helping people, that's not good enough. You need to be helping people more than if you just gave them the money that it costs to help them. So I think they do a lot of cool things and I think they've been very successful. I think they do... They're probably going to downsize a little bit this year, just because they have been raising a lot of money and I think it's not going to be quite as much this year. So definitely plenty of room for more funding there.

08:16 HK: Next set of charities is deworming charities. So there's a set of parasitic infections that are extremely common, especially in Africa, also in India. These are the roundworm, whipworm, hookworm, and then this worm called schistosomiasis. What we know is that a large, large number of schoolchildren are infected with these worms. What we don't really know is how that affects them. There's very little in the way of solid, clinical symptoms of having a worm infection. If you get a severe infection, you start to show symptoms, and you generally go to a clinic and you can easily get treated for it. The question is what do you do with people who have mild or ambiguous or not clearly symptomatic infections? And deworming is the process of going around and basically treating children with drugs once a year generally, or something like that, that kill the worms. And so you treat everyone whether or not they have worms. You don't bother to test them because that brings the costs way up. The drugs have very few side effects.

09:20 HK: And then the question is, what does this do? Well, it definitely kills the worms. Does this help people in any measurable way? And that is the subject of a massive debate. And so if you Google "Worm Wars," you'll see economists getting very angry at each other. A lot of people getting angry at us. Us hopefully not appearing angry when we write about it. [chuckle] The bottom line or how I would summarize the worm wars really quickly, and there's a lot more depth to it, there's a lot of short-term studies and they really can't pick up much of anything, in terms of weight gain, in terms of better iron content in the blood, in terms of anything with test scores. Basically in aggregate you see almost nothing from short-term studies of deworming.

10:01 HK: And then there are long-term studies, very, very few studies. And the studies have significant flaws, significant problems, and they find amazing effects. And so how you put together all that evidence and decide how you feel about deworming is something that we have put a lot of effort into. Certainly in the past we've picked apart a lot of

studies, we've gotten their raw data, we've gone through it ourselves, we've looked for all the issues. I think one issue we've really harped on that I think doesn't get enough attention is that the best study on long-term impacts of deworming, the deworming took place in one of the most worm-heavy areas, and it was in the middle of the El Nino, where there was a massive amount of flooding, and worm infections went to totally unrepresentative levels. And so anyone extrapolating from that study I think needs to be more careful than they usually are.

10:50 HK: So we definitely have lots of reservations. This year we got even deeper into it. David Roodman, who works for us is kind of... He's a person who wrote an entire book about the evidence behind microfinance and he tends to go very deep. And he took the deworming studies, got the data and just replicated everything himself, subjected it to a million sensitivity analyses. The bottom line is one of the studies that was encouraging about deworming we no longer believe. Another study we think has still got a lot of problems, it's the El Nino one, but it has stood up to his scrutiny, which I think is pretty impressive for a study. And so our position on deworming is basically the same as what it's been, which is that, I would probably guess it's more likely than not that deworming actually doesn't have any substantial impact on people's lives, because we're just working with a really thin evidence base. But, if it has the kind of effect that you've seen in those studies it's an unbelievable deal for your money. It's 50 cents for a dollar per kid per year to deworm someone.

11:52 HK: And so when you net all that out and you look at it, we end up thinking that deworming is about about a five to 10 X better deal than cash transfers, in terms of how much good you're accomplishing per dollar, and you can measure that in terms of, you deworm someone, they make more money later in life. How does that additional money compare to the additional money you could've just given them, that they could've then invested? I'll take questions in a minute. Our take is that when you net all that out, deworming looks better than cash transfers, maybe five to 10 X, and looks about as good as the malaria interventions that I'll get to. So then in terms of the organizations that do deworming, there's Deworm the World... There's four of them, so two of them are new this year, or new to our list.

12:36 HK: There's Deworm the World, which is our favorite in terms of the cost-effectiveness, which is generally a function of how many worms there are where they're working, and our favorite in terms of their monitoring reporting. But they're the ones with the least financial need. Basically it's not clear that they need any money and if they do it's not a huge amount. And then there's Schistosomiasis Control Initiative, a group we've recommended for a long time. Historically, we've really struggled with them. We've really had a lot of trouble communicating with them, not felt the quality of their reporting was high. We think they've been improving, so I wouldn't say, I wouldn't say A+ still, but I think they've improved a lot, and we now think they're second out of the deworming charities in terms of our confidence, and they have a lot of room for more funding.

13:22 HK: Then there's Sightsavers, which is new to the list. They're a group that does a whole bunch of programs that are focused on preventable blindness and vision issues. They have a new deworming program and based on their past non-deworming programs that are similar, we feel confident enough in them to recommend them, but we're going to need to follow them for longer before we really understand everything they're doing as well as we do the first two. And then the END Fund, I would say is a similar situation where they do a bunch of programs with neglected tropical diseases, some of that is deworming. We think they're good enough to recommend in terms of the evidence we've seen, but it's going to be a couple of years before we have the same level of confidence in them.

14:05 HK: And so then the final category of charities I'll talk about is malaria prevention. One of our top charities for a long time has been the Against Malaria Foundation. They distribute bed nets with insecticide on them and those hopefully block and kill mosquitoes and prevent people from getting malaria. The main benefit of this is preventing deaths of children under five. If that's something you're excited about, you should be excited about this. If you don't value that, as some people have conflicting opinions on this, you may want to give to another charity. There are other benefits for preventing malaria, but that is the main one. AMF is a good organization, we've been working with them for a long time. They have better reporting and monitoring than by far any other bed net organization. This year they expanded a lot, they signed a lot of new deals. We got a little bit less confident in their monitoring. They have a massive funding gap. So in theory they could use something like \$80 million plus. I do think the more they get, the more the cost-effectiveness falls a bit, because you start getting into the possibility that they're pushing out other bed net donors, like governments, and especially because they have a preference for doing really big distributions, they'd have a preference to do a whole country. But we still think they're excellent.

15:20 HK: And then there's another group, Malaria Consortium, which does not do bed nets. They do seasonal malaria chemoprevention and that is, it's almost like a hybrid of bed nets and deworming. It's drugs to prevent malaria for young children under five. It's like if you travel and you take drugs to prevent malaria, this is that same basic thing, aimed at

people who are especially vulnerable. So Malaria Consortium, actually kind of a similar... That thing we think the evidence base is not quite as good as bed nets, but it's kind of similar. The cost-effectiveness is not quite as good, but it's in the ballpark. And over all, these malaria groups, we would say the evidence is stronger than for deworming, it's not as robust as cash transfers, and we think the cost-effectiveness is maybe a little worse than deworming, but very close, and a lot better than cash transfers.

16:10 HK: So that is our seven top charities. We also recommend a number... Not recommend, we list a number of other charities that are noteworthy in various ways on our website, but I've been going for a while and I'm not going to get into that. The two big takeaways to just note are, [1] we put more time and we were able to do more on finding new top charities this year than in the past, and we think they have a lot of room for funding. And also, Good Ventures, which is the major donor we work with, is giving \$50 million total, but they're not giving more. And I can get into why that is. I will certainly talk a lot about that at the next session, but that leaves a really massive amount of funding gap. And so there's really a lot to give this year. There's no practical limit to how much they can absorb relative to how much they have room for. And so that's kind of where we're at, and those are the big updates this year. And with that, I'll take questions for a bit.

[background conversation]

17:11 HK: Oh, and I'm going to repeat the questions for the recording.

17:15 Speaker 2: When you said deworming is five to 10 times more effective, is that contingent on thinking it's effective, or that's adjusted for...

17:22 HK: That's adjusted. So when I say deworming looks five to 10 X better than cash, that is adjusted for the fact that we have a lot of skepticism about the evidence. So if we took the evidence at face value, it might be something like 100X or more.

17:39 Speaker 3: Malaria Consortium. I know it does a lot of things beyond chemo. It's like a general malaria advocacy kind of organization that works with many strains of health. How would you... Are you recommending a restricted donation to them?

18:00 HK: Yeah.

18:00 S3: What's this link between this relatively uncommon intervention and a big organizations that does a lot of things.

18:07 HK: Yeah, Malaria Consortium is a big organization that does a lot of things. The same is true of all three of the new top charities. So Sightsavers and END Fund, none of these are focused entirely on the program we're recommending. In all cases we're recommending donations that are earmarked for that program. And we do believe that those donations will cause that program to expand more than it would otherwise, that's definitely something we've checked out. I should make one more comment about the top charities. Just a good thing to just know on the progress on the GiveWell front, I was basically not involved this year in any of the research, kind of came in at the end and played with the cost-effectiveness file, but didn't really have anything to add. Basically everything that I gave you just now is mostly from reading our website, or basically all of it.

[laughter]

18:54 HK: Before I came here I was just on our website making notes. So I think that is a good step, that is good progress for GiveWell as an organization, that I'm definitely unnecessary now. [laughter] And also another reason for you to know that is that there will be a lot of questions that stump me tonight. So that's a warning, yeah.

19:12 Speaker 4: Is the expansion in the number of top charities something that you expect will continue?

19:16 HK: Is the expansion in the number of top charities something I expect will continue. We're definitely not going to double our list every year for very long. But I do expect new top charities next year, not for sure, but that's something I'm going to get to. I'm going to take a little interlude and talk about that. But in addition now to recommending top charities, we do a fair amount of work trying to help facilitate the emergence of new top charities. And we're starting to call this the incubation program. We used to call it GiveWell experimental. And so there's a bunch of groups right now

that actually we have recommended grants to them, from the major donor we work with, that are trying to help them cross along that path to get to a top charity. Some of that is for research, some of that is for charities themselves, and there are a couple that I think are reasonably likely to get on that list either next year or the year after. I don't think there'll be three new ones next year, but there might be. Yep.

20:10 Speaker 5: Have you ever or would you ever look not just at donation of the money, but actually donation of time and how could that be the most effective way?

20:18 HK: Sure. Do we do look at donations... Do we ever look at donations of time instead of money? So GiveWell is just focused, it's got its mission and the mission is money. The reason we're more focused on money is because we feel that we can have more impact that way. So I think one of the issues that we do see sometimes with donation of time, A, it's hard to give general advice on it. Everyone has different skills and so it's harder to get a scalable recommendation. I will also observe and we've written about this a couple times, to do really good work for an organization, a lot of times you need to be trained, you need to have context, you need to have experience, and so it can be hard to get a lot of real impact out of someone who's just donating a little bit of time. Especially if they're not... If they're changing their mind about how committed they are, and if they're not always into it, and that's just something to keep in mind.

21:12 HK: I actually think sometimes when the charity takes on a volunteer they're really trying to recruit a donor or a fan more than they are trying to get a volunteer. And usually in the past... We don't even take volunteers anymore. In the past when we've taken volunteers it's generally just been for recruiting, in the hopes that people turn into employees. Those are some general observations on that. I think it's very challenging to figure out how to do a lot of good when you're donating time. And I often encourage... It depends who you are. I often encourage people to just not do it, and focus on doing things that they have a huge comparative advantage in. But it's not something we're experts on, and it's just not where we really focus. Yeah.

21:54 Speaker 6: So my understanding is GiveWell's previous cost-effectiveness models... So the benefit of saving a child's life was considered 36 quality-adjusted life years which was supposed to represent basically a whole life, with some discounting. And now it's been reduced to nine years. So, I guess two questions, one is, what factors went into that reduction? And also does it even... If you're no longer saying that a child not dying is worth a full life's worth of quality-adjusted life years, why even use quality-adjusted life years, isn't that how it [22:39] \_\_\_\_ Why not convert it to dollars? This is what it's worth in dollars.

22:44 HK: Sure. So this gets into our cost-effectiveness estimates. This is an observation that, when we're estimating, when I give you these numbers, like X is five times as good as Y. To get to that number we need to build a model and the model needs to do things like, well, AMF is saving lives of children under five. Schistosomiasis Control Initiative is helping people earn more 20 years later. How do you even compare those two things? And our answer is subjectively. So we have we have a spreadsheet and the spreadsheet asks you a bunch of questions. And some of the questions it asks you are like, "How many years of healthy life do you think are equivalent to one life saved of a child under five?"

23:27 HK: Because some people might think, "Well, that's a whole life lost," and so that would be 30-plus years, and other people would think, "Well, I don't really think that children of a certain age, maybe under one, have an identity yet. So I don't want to count it the same way, but I do want to account for the suffering of families and things like that." And so people can debate that all day. But methodologically, and they actually do at GiveWell, but methodologically the high level answer to your question is that we have a spreadsheet where everyone can confront these questions, put in their own values, and the spreadsheet will spit out what the upshot is in terms of your values, how you should value the charities compared to each other. When I give you...

24:04 HK: So anyone can do that, anyone can go to our website, download that spreadsheet, put in your own opinions, figure out where you stand. When I give you a number that's like 5X, where that's coming from, is that all of our staff members who wanted to, including myself, went into that spreadsheet, put their thing in, and then we took the median. And so for a particular value... I think that the nine DALY's figure you're talking about, is probably a median of staff. And so that's probably like some people think that it should count for a whole life, and some people think it should count for nothing, and some people are really confused, and they put a number in between. They do it on purpose, they're saying, "I'm confused, I'm putting nine." So, that is methodologically how that works. At the same time, the differences are stark enough that, you don't have to trust us completely, and you don't have to download a spreadsheet. You can think about how you value lives of children under five, and if it's something that, basically to you seems totally inferior to raising incomes, then I think I've made it pretty clear which way you would go with that. Yeah.

25:04 Speaker 7: I have two questions, one was about earmarking. I have a sense that before you guys thought that [25:10] \_\_\_\_\_ wasn't a good idea. Is there a reason that in this case that I should believe they're not going to just reallocate resources?

25:16 HK: Yeah.

25:16 S7: And another question is about GiveDirectly specifically. I always thought that a big part of their value add was the fact they're doing conditional cash transfer experiments, and they're trying to influence NGO governmental policy. You didn't touch very much on that. Do you not feel it's a large part of the value add?

25:36 HK: Sure. Okay, so first on earmarking. Yeah, we've definitely pointed out in the past that... I think it can definitely be dangerous to just... Let's say you have a charity that does a whole bunch of different things, and you like one of the things. So, you have a charity, let's say, that do vaccines and job training programs, and you like the vaccines, and you don't like the job training programs. If you write them a check, and you say, "Only use this for vaccines," I think the default outcome of that, and a thing that's going to happen a lot, is your money won't really have that impact. They will... Let's say that I'm the charity, and next year I want to spend 10 million on each. So now I'll take your money and I'll put it in the vaccines pile, and then I'll take other money that I have flexible, and I'll shift it over to the job training pile. So, that I think is something that you have to watch out for.

26:24 HK: Why are we recommending earmarking this year? I don't actually have an answer. This is something I just don't know, except that I believe we watched out for it, [laughter] and questioned the charities enough. Certainly, there are things you can do, you can just say, "How many treatments are you going to deliver next year if you don't get any more funding of any kind? How many are you going to get if you get this much restricted funding?" And then we can hold them to that, based on how much comes in. So that's what I'm guessing the basic MO is. Also, earmarking is less and less of a problem the more and more you're giving. And so when we're dealing with aggregate numbers tracked by GiveWell, they're going to be in the millions. That's a much harder thing to fudge. I think if someone gives \$100 earmarked, that's just very likely to be fungible like that. So those are all guesses, 'cause this is not something I got into the details of. I think it's pretty straightforward to verify, it's just something you have to verify.

27:18 HK: On the second question, is a big part of GiveDirectly's value, the way that it's influencing the policy conversation, the way that it's making people think, maybe we should hold charity to a higher standard? I would say it's definitely a part of their value. I would say they get a good amount of buzz in the development economics community. And like I said, they have a partnership that I don't think is fully announced yet, but it's with a major institutional funder to explicitly take some common aid program, and compare it to cash transfers. So, I do think GiveDirectly has an impact on the dialogue.

27:52 HK: The things that I emphasize when making the case for these charities, is the things that are the least subjective, and the easiest to make the case for. And I think, especially GiveDirectly, it's already doing pretty well, and it's already pretty big. So if you give them an extra \$100, probably the main thing to focus on, is that 82 of those dollars are going to someone who really needs them. But yes, I do think there are these other kind of intangible effects of GiveDirectly. And I think GiveDirectly is still, even on the margin, I think GiveDirectly is stronger on that than the other top charities. I think they are just way more aggressive about looking for creative, ambitious things to do and ideas to spread, and I think that is an advantage of GiveDirectly. If you like that, it's a reason to give there. Yep.

28:34 Speaker 8: Have you guys looked into the cataract surgery charities, because you mentioned the Sightsavers? So in terms of quality-adjusted life years, does that make sense?

28:44 HK: Yeah, we looked into cataract surgeries. So cataract surgery, basically it's a very cheap, simple surgery that can restore someone's vision or prevent them from going blind, and that seems like it could be very cost-effective. I know in the past we've looked at cataract surgeries, and we've just run into the earmarking room for more funding problem, where it was kind of like, there's a charity, they do a bunch of surgeries, we say, "If you get more money, will you do more surgeries?", and they say, "No." But, I do think it's possible. I think we're still interested in that, and we're still looking for a way to turn dollars into cataract surgeries. And I think that, I would have to check, but I think that may be a category that is possible for a future top charity. In terms of there may be... Some of these top charity investigations take years, and so there may be groups that are being investigated right now that do that, but I'm not sure. Kat, do you happen to know?

29:37 Speaker 9: I know we published a new report on cataract surgeries this year, and I believe room for more funding

was the issue which you mentioned. So, your understanding is about what mine is, that there's a new report out there that came out in August that they published.

29:50 HK: Cool. Yep.

29:53 Speaker 10: How much influence would you say that GiveWell has on the organizations that you guys promote? So, are you viewed as a big enough source of income that you can even say, "We think you guys would really benefit from doing this study to test this factor"?

30:07 HK: Sure. How much influence do we have on the organizations we recommend? I think just as a first comment, I think we probably wouldn't ask an organization to do a study, because we think that usually that's done better done by some other organization or by academics. We might ask them to partner on something. I would say that in theory, we could be pretty influential. I think we also have to be careful with how we use that influence. We don't want to turn this into a bunch of GiveWell-run organizations. I don't think that would make them better, and I think that would make our recommendations a lot less credible, so when we are making recommendations, they're usually just ways to... I think the main thing we do is we tell everyone how they're scoring on our criteria, and why they're scoring that way, and I think that definitely can influence organizations to do things differently. So, we've told SCI for years, "You're getting a lukewarm recommendation in a sense, because of the communications issues we're having and because it's difficult to track where the money's going and it's difficult to understand what you're doing." I think there's been a lot of improvement on that front. How causal was that, I don't really know, but I feel fine if it is.

31:17 HK: And then, on AMF, we've certainly made it clear, for example, we said... At one point we suspended their recommendation, because we didn't feel that they were going to push more money into the field, and we made it clear that we were happy for them to make certain kinds of compromises that we didn't think were going to hurt the cost-effectiveness of the program in order to get more money spent, in order to get more bed nets delivered. Again, they did end up moving in that direction. Again, in both of these cases, I don't think it's clearly causal, I think that these are both cases where it could've been totally independent of us, but it's generally "Here's how you're scoring on the three criteria, here's why." I think with some of our incubation charities, it can even be a matter of, like, "This program is not very cost-effective, here's one that is," and they do what they want with that information and certainly there's a reason to take it into account.

32:15 Speaker 11: Are you looking at cost-effectiveness and growth, like kind of immediately you could say AMF spend X million dollars that benefit now and saved some lives or they can spend that same money for a marketing budget and increase donations next year?

32:32 HK: Sure. Yeah. How do we look at the question of should AMF spend the money delivering bed nets or should they spend the money on fundraising and then that causes them to bring in more bed nets? Yeah, I think it's fairly rare, I would say, for fund-raising to be a big enough part of an organization's budget that it would have a huge negative impact on our cost-effectiveness estimates. There's not charities that are getting like knocked off the list, 'cause they spend too much on fund-raising. Anyone who's getting this far, generally that's not the issue, if they're doing an effective program, and things like that. So I would say it mostly hasn't come up. I mean, we're also, on the flip side of that, we're not going to recommend a charity whose main value add is fund raising, that's like a different kind of value proposition, and that's not the thing that we offer as GiveWell, so yeah, that's basically that. I'm going to also talk now just for a couple of minutes before I take the final batch of questions, about that incubation stuff that I mentioned. So I thought this might just be a good time to mention a little bit more specifics on what we're doing there.

33:34 HK: So like I said, we're very interested in facilitating the growth of future top charities. We work closely with Good Ventures, which is a major funder. They're interested in the development of new top charities. And so, examples of some things we've done, GiveWell Incubation is not the same, it's different criteria. So GiveWell, you need to be proven cost effective and scalable. For GiveWell Incubation, all that needs to be the case is that we believe that there's a decent chance you're going to turn into a top charity and that this money will help you, and that it's a good deal. And so, some grants... For that reason, we use actually a process more similar to the Open Philanthropy process, which I'll get to. But things we've done there, there's a group called New Incentives. They do conditional cash transfers. So those are cash transfers that people have to do or prove something to get. And an example of what they were doing is, they were trying to get HIV positive mothers a payment in exchange for provably taking this preventive therapy that prevents HIV from transferring to the child. And so, you're kind of get a two for one, because you're getting money to low income people and you're also, if you pick the right outcomes, you're getting activities that might help benefit their children.



34:47 HK: That was a group that we supported for a while. We were happy and impressed with how they were scaling up and executing on their program, but in terms of enrolling HIV positive mothers, we didn't feel that the... The numbers were not what we hoped they would be, and therefore the cost-effectiveness of the program was not good enough to recommend in this time around. And so now, they have been talking about going and doing a different kind of conditional cash transfer, more related to immunization, getting children immunized, and that one scored higher. And so, we're continuing to support them as they try to find a way they can do something that's going to qualify to be a GiveWell top charity.

35:26 HK: Another group we've supported this way... So we supported Evidence Action, which runs Deworming the World initiative. We've also supported the piloting of this program, No Lean Season. And that is a program where they give people vouchers to go from rural to urban areas in the developing world, and a lot of times what that means is that people find they can earn more money, and then a lot of times in the studies, many people went back to the city the next year without even having a voucher. So that's a good sign that this is actually benefiting them. There are some good studies on this, some good evidence, and we've been funding pilots in the hopes that they will get to the place where they can deliver a lot of it, and we feel like it's going well, and we're ready to recommend them as a top charity. That may happen in a future year.

36:08 HK: Supported a group that works on pneumonia treatment, it's about \$6 million for initial scale up, and scale up with monitoring to see how this goes at scale. And then they are a contender, they may be a top charity next year or the year after. There's another group, Charity Science Health, that was explicitly founded by some people who said, "We want to be a GiveWell top charity. Our mission of our organization is to be a GiveWell top charity." They're working on SMS reminders to get people immunized more reliably, and that could be quite cost-effective. So that's a group that we have done some supporting of in the hopes that they will get there.

36:43 HK: And then there's IDInsight is an example of where we've supported something more along of the lines of research and monitoring. So this is a group that, they basically help collect data and evaluate programs in the field, and unlike a lot of the other groups that do this... So a lot of the groups that are known for doing studies, they're academic groups, and so a lot of times their interests are different than ours. A lot of times they're more interested in illustrating important academic concepts or in challenging important academic concepts, and less interested, in let's say, doing the fifth study on something to make sure that it really works the way we think it did, or I'm just seeing whether a program is being delivered the way it's meant to be delivered.

37:25 HK: IDInsight is an evidenced-based, study-based group that does more of that. We've given them general support, and we've been working with them on trying to figure out where they can go and collect better and more data on GiveWell top charities, or potential future GiveWell top charities, that can help us make better decisions. So that's the incubation work we've been doing. And I also have not been very involved in that, but that's some of the things we do for future top charities. Yeah.

37:52 Speaker 12: So two questions about those. The first is about Charity Science Health, and whether or not you think with the strengths and weaknesses their approach were for [38:01] \_\_\_\_\_. And then on IDInsight, I believe they're a for-profit, I could be wrong, but I think they certainly have clients, and just how you handle that combination?

38:13 HK: Sure. So first on Charity Science Health, unfortunately I don't have much to say about the strengths and weaknesses of their program. I wasn't involved in that intervention report, and would have to just go on the website and see it. But we've certainly written it up. And then on IDInsight... What was your question about IDInsight?

38:32 HK: I think that they're for-profit, or that they at least have paying clients that they're funded on.

38:35 HK: Yeah. I don't think they're for-profit, and I could be wrong about that. But I think they do have paying clients. So they'll go to a charity, or a government or something, and someone will pay them to figure a specific thing out. Yeah, I guess... I'm pretty... I would be a little surprised, 'cause I feel like this would have come up as a logistical challenge if they were for-profit. But it wouldn't necessarily rule it out, it would mean that the money had to be accounted for in a certain way.

39:01 HK: The other client necessarily explicitly figured out that they do take money from clients, if that's [39:07] \_\_\_\_\_?

39:09 HK: We fully... So, have we looked at the fact they've taken money from other clients? Yeah. With any

organization we make a grant to, we're always asking the question, "Alright, what's all your revenue? Where's it coming from? What's all your expenses? Where are they all going to? If nothing changes next year, what are you going to do? If you get more money, what are you going to do?" And then we can hold them to that, the amount of money they actually end up getting. And so that's always the basis for our recommendation. So yeah, my guess that IDInsight, is that they have a lot of specific clients wanting them to evaluate a specific program, we want them to look at top charities, and potential future top charities, and believe they need more money to do that, also believe that they could use more general capacity to just try things. Yep.

39:52 Speaker 13: A lot of the top interventions seem to be health-based. And could you speak to why that might be on a very broad level?

40:00 HK: Yeah, why are so many of our top charities health based? It's definitely true, I think the nature of our criteria creates a bias in some ways toward the measurable, and health tends to be more measurable than other things. I think some other qualities that... And again, this is a bias toward the measurable, it's fully acknowledged. I think it is quite logical for someone who isn't spending all their time engaging with questions about where to give, to give to something where there's a case that can be made, and that involves things that are more measurable. So in some ways, it's like the idea of dropping your keys, and looking first under the lamp post. People often present this as kind of a joke or a fallacy, it's often a rational thing to do. If your keys are sitting in a place where you can see them, that is going to be the quickest way you find them.

40:50 HK: And so I think it is actually quite rational to have a bias toward the measurable, when you're trying to make a tough decision, and you don't have the context and the soft evidence to do something a lot riskier. Other stuff about health. I think, that said, I don't think we're sacrificing necessarily much cost-effectiveness by going for the measurable, because I think there's a relationship where health charities, they tend to be doing things that... They tend to be helping people in ways that are very universal, and very robust, and very generalizable and scalable. And also because we understand things so well, and because we're able to study them, you're often able to strip out... You're often able to take the way you're helping someone, and strip it down to its essence, and get it as cheap as possible, and distribute it as widely as possible.

41:36 HK: So I could do a job training program, or I could give out bed nets. But with the bed nets, we've isolated what to do that helps people, down to the point where it costs less than \$5 a person. So I actually... It's not the case that there's a bunch of global health or... It's not the case there's a bunch of global poverty charities that I think are more cost-effective, and they just happen to not meet the criteria. I actually think that our top charities are doing things that are incredibly cheap, partly by virtue of their measurability and understandability, that are effective and you're getting incredible value for the money, and I think that's a common feature of health programs.

42:10 HK: I think it's also just kind of the case, and we've written about this a couple of times in the blog, but we are trying to help people who are far away who we don't know much about. And I think helping with health is something that donors and governments and well-meaning western world has historically done pretty successfully, a lot of other things they've historically done really badly. Health is a pretty universal thing, it's a thing that we understand, it's a thing we think we can help with, and I think it is quite appropriate when trying to help someone that you don't really understand, it's quite appropriate to focus on effective things you can do to help. So that's some of the health stuff. Yep.

42:49 Speaker 14: Are you optimistic about future technologies such as CRISPR? I don't know a ton about CRISPR, but I'm hoping that there are other technologies that are potential future interventions for things besides malaria.

43:05 HK: Sure. Am I optimistic about future technologies? I think you're referring to future technologies that might help find other ways of controlling these diseases we're fighting. So for example, there's been some press about... And I think there was... There is a bit of material on our website if you Google for it, these gene drives, a gene drive being like basically a way of intervening in a mosquito population so that you stop the spread of malaria. And that's using genetic engineering techniques, and I'm optimistic. I definitely think there's a chance that some kind of alternative new tech like this will turn out to be a much better way of fighting these diseases than bed nets. I will say that we believe... We have made at least one grant along these lines. We are in favor of it. We are looking for opportunities to help with it, where it seems like a good deal and where we can help, and if at any point, we believe that bed nets are no longer needed, because it's better to do some other intervention to control the mosquito population, then we will adjust our recommendations accordingly. Yeah.

44:17 HK: Where does GiveWell stand in terms of outreach, how are you all spreading awareness about GiveWell,

because I know if we could get more donors to be looking at GiveWell recommendations and [44:32] \_\_\_\_ charity evaluator recommendations, would it have a more robust influence?

44:38 HK: Yeah. How is GiveWell thinking about outreach and what do we do to spread the word about our own work? To date, I've kind of said the same thing whenever people asked me this question for a long time and so it may be... I don't know, may be frustrating people or something, but for a long time, we've just said, "Our priority is not outreach, our priority is research", because we have to have the good research, and we're still seeing good growth and our money moved. So we track the donations people give to our top charities and that number has grown pretty strongly every year, we'll see what happens this year. In the last year, including Good Ventures, it was over \$100 million and about 70 of that was Good Ventures and about 40 million was coming from other people. And actually, for several years, it's been a little bit challenging to even find enough recommendations to take all that money. This year that's not the case, this year there's plenty of room for funding. But yeah, I would say that we've really focused on getting the research done well, and also on getting the research to the point where it doesn't need the founders as much anymore. And that's been our top priority and so, outreach has been a secondary thing.

45:43 HK: And there also have been other groups that have done their own outreach effort. So especially groups of the Effective Altruist community, there's multiple of them that are specifically trying to... Their mission is to get people to give more to effective charities, which could mean top charities or something else. That's what they focus on. What we focus on is just getting the recommendations right.

46:05 HK: So, I would say that in the future, if we see our money moved slowing, and we have lots of room for more donations and we're not getting a lot more donations, and we have the organization pretty well under control, where we're able to do quality research without too much reliance on the founders, that may be a time to really ramp up outreach. For now, it's more of a low-hanging fruit thing. So we do the online advertising, a very important part of our outreach is talking to donors and just building relationships with them, getting them more confident in our work. We do give talks, but it's not the main thing that we work on. I take one more question, then we're going to take a break. Yep.

46:46 Speaker 15: Does GiveWell look at programs that are designed to not to need cyclical funding, like social businesses, once you fund them, they become self-sustaining by design?

46:56 HK: Sure. Does GiveWell look at giving opportunities that can become self-sustaining? We did this for a while at one point. So, we spent quite a bit of time looking at micro finance, for example. In the time we spent on it, we didn't end up feeling that there's anything really repeatable, scalable proven there. I think there are definitely times when a small investment upfront can create something that then sustains itself and grows, and a lot of times those look like for-profit investing opportunities, and that's great and that's exciting, but I can't point you to any kind of thing where it's like, every time you give a million dollars here, you're going to get a self-sustaining thing. And furthermore, I didn't even see anything that looked like I could point to it and say, this is a better deal than top charities, those in that category. Because I think, for example, micro-finance, I think the impacts on poverty and the impacts on people's lives are just super ambiguous. So, that's the story there, but it comes back to those criteria. Alright. Well, we're going to take a break of some amount of time.

48:02 S?: 10 minutes.

48:02 HK: 10 minutes and then we're going to talk about Open Philanthropy Project. So thanks, everyone.