

## GiveWell San Francisco Research Event, June 14, 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.*

0:00:03 Elie Hassenfeld: Hey, everyone. Thanks for coming. I'm Elie Hassenfeld, I'm one of GiveWell's co-founders. I'm gonna be talking to you, for the next hour or so, about GiveWell. And then Holden, who's one of the other GiveWell co-founders, is going to talk about the Open Philanthropy Project in the second hour.

0:00:21 EH: I'm gonna primarily cover two major topics about GiveWell. One is going to be an update on what we've learned on our top charities during the first half of the year. I'm gonna focus mostly on the Against Malaria Foundation, where we've had the most important updates. And then also I'll pause, take some questions on that, and then talk about work that we call GiveWell Experimental, which is supporting the development of future GiveWell top charities. These updates are definitely going to assume some level of background knowledge on GiveWell. Looking around at the folks who are here, I think a lot of you have a significant amount of context on what we do. But during the Q&A, don't hesitate to ask basic questions if you think you're missing important context. Please go ahead. We're experimenting with different approaches to these events and still trying to find the right balance between basic information and detailed updates.

0:01:18 EH: And then finally, as you probably expect, we're recording this event, and we're planning to publish our recording in transcript on our website. We are happy to cut anything from the recording that one of you want. If you ask a question and then after the fact, would prefer that we not include your question online. Feel free to come and let me know. Send me an email or send an email to [info@givewell.org](mailto:info@givewell.org) after the event. These are not going to be made public in the next 48 hours, so you have some time. If you're concerned about anything and want to let us know to not put something up.

0:01:56 EH: So in terms of updates on our top charities, I think the place where we have the biggest update so far this year relates to AMF, the Against Malaria Foundation. They support distribution of malaria nets that strong evidence shows save children's lives from malaria in sub-Saharan Africa. They were our top recommendation to donors at the end of 2015. And due to our recommendation, they received a significant amount of funding on the order of \$40 million. All the money that we directed to them and the money they have on hand, they have spent a significant amount of it, leaving them with roughly \$20 million on hand that they can allocate to future distributions.

0:02:45 EH: Now, one of the big questions that we had about AMF at the end of last year was whether they would be able to successfully sign new distribution deals with partners. Now they've done that. So, we think this is a significant positive update for our view of AMF and its ability to successfully take in money and use it for bednet distributions in the future. And at this point, given the funding opportunities available to them, we think that even beyond the \$20 million they hold, they will be pretty likely to be able to commit another \$25 Million beyond that. And so, given the amount of funding that we expect to be able to influence in the next six months, we think AMF will be in a situation where their work is significantly slowed down due to lack of additional funding.

0:03:42 EH: Now, one of the things that we've been talking about recently, internally, is whether we should recommend that donors move gifts up, that they would have otherwise expected to give in December when the vast majority of individual donors give, to give in June or July to enable AMF to move forward more quickly. This is something that I'm personally thinking about, whether or not I should do. We'll make a final decision in the next couple of weeks and we'll publish our recommendation to donors. We've also been talking about whether or not we should recommend that Good Ventures, which is a foundation that we work very closely with... Carrie is the president of Good Ventures, she's here... Whether we should recommend that Good Ventures make an additional grant in the nearer term to AMF.

0:04:32 EH: Now, we are currently not expecting that we're going to recommend that Good Ventures will make an additional grant. The reasoning behind this is something we talked a lot about at this event last December. We also wrote a lot about it. We called this Coordination Theory, which is basically the question of whether a donor that has the capacity to fill a gap should fill the full gap, or should explicitly hold back to try and draw in other donors. And so for the time being, we're still planning to recommend that Good Ventures hold back in June and reconsider what they should do at the end of the year, or what we should recommend that they do at the end of the year, as part of our standard recommendations to them about their end-of-year giving to top charities. This is probably one of the subjects of debate that we've argued about most intensely, internally, over the past couple of weeks. So definitely, I'm happy to take any questions on that topic.

0:05:29 EH: I think the biggest substantive of update about AMF though is that, given their success this year signing deals for tens of millions of dollars in multiple places, we have significantly more confidence than we did six months ago that they'll be able to take in additional funds and utilize it to fund bednet distributions in sub-Saharan Africa.

0:05:54 EH: One of the biggest changes in how we've thought about what we call room for more funding for AMF, which is how much funding we think they could effectively absorb and then distribute out, is that in the past, we had relied heavily on AMF's own internal estimates of which countries and distributions it was likely to negotiate with in the near future. And the biggest change this year is that, because of our increased confidence in AMF, we think that virtually all of the global net gap, meaning all of the places where people are not covered by nets and other funders are not coming in to fund them with the exception of, one, places that are so difficult to work that we doubt that AMF would be likely to work there, or places that the needs are in aggregate so small, that AMF would be unlikely to prioritize working there. With those two exceptions, the gap that we see for AMF is driven by the global need for bed nets, through the end of 2017 and we estimate this total gap at around \$40 million above what AMF currently has on hand. This is the gap that we expect AMF to struggle, to not have funding, to try and fill over the next several months and it's the gap that we'll be seriously considering as we make our recommendations at the end of December.

0:07:32 EH: The biggest questions that we still have about AMF, and that are counterweight to the positive things that I've said so far about our increased confidence, there's basically three of them. So first, AMF has an improved track record of signing additional distributions, but they still have a very short track record of distributions for which they have collected and published monitoring on how well the distribution went.

0:08:02 EH: So the monitoring that we have where organizations they've partnered with have gone back and checked, whether or not nets were owned by households that AMF targeted or AMF's partners targeted, that they were hung directly in those homes. But still we only have that data and that data's only been published for one country and one distribution partner of the roughly four or five deals that AMF has signed, because many of them are ones they have just signed this year and are not ones that they could yet have data for.

0:08:34 EH: And I think a major question that's worth keeping in mind, is whether the track record that AMF has is sufficiently strong to warrant additional funding to AMF at this point. We think it is, but I think it's certainly an open question. The second major question for bednets as a whole is the question of insecticide resistance. When I say malaria nets, these are mosquito nets that are treated with insecticide, and they not only provide a physical barrier against mosquitoes, but the insecticide repels and can kill mosquitoes. Over the last several years, researchers in the malaria community have become increasingly worried about the development of resistance among mosquitoes where they're no longer as susceptible to the insecticides on the nets as they previously were. Now, we know for a fact that mosquitoes are developing resistance, based on lab studies where mosquitoes are put into contact with insecticide.

0:09:34 EH: We have very little information about the degree to which this leads to significant decreases in net effectiveness in the field and our best guess now is that it reduces the effectiveness of nets by something in the order of 20% to 30%, that is our best guess based on pretty limited information. There are several studies underway, that are testing in a rigorous way bednets in the field in insecticide-prone areas, and we expect to have some material updates to this within the next 12 months or so.

0:10:14 EH: The final question that we are thinking about with AMF and it's another one that we dug into the first six months of this year is the question of whether in instances where AMF provides funding to a country for its bed nets distribution, what's the counterfactual? What would have happened had AMF not provided that funding? And the way that we tried to look at this is we looked back at five cases since 2013 where AMF had planned or had engaged in negotiations with a country and then ultimately did not follow through on providing funding to that country for one reason or another. And in those cases, the weighted average of what happened, weighted by the size of the distribution, is that roughly speaking, those distributions did end up being funded about 16 months after they otherwise would have been. And the expectation is that a bednet will last about two and a half years or a little more than two and a half years, and our best understanding is that the distributions that AMF did not fund in the past or AMF funded in the past, I think a reasonable guess is that, they would have been filled about 60% of the way through the time for which it would have been a sort of total impact by AMF where they had fully filled the gap that would not have otherwise been filled.

0:11:43 EH: I think there are a lot of questions about what you should do with this information. One question is, what the funds that came in would have otherwise gone to? I think that it's reasonable to argue that they would have, otherwise gone to fill some other bednet gap in some other location, or they would have gone to some other plausibly

effective activity. These are other questions we've been debating and don't have significant information about. Another question is how representative the past four years are of what we would expect the future case for AMF filling the funding gap to be, and I think if anything, are not particularly well informed with information guess is that it's more likely that other funders will have more funding available for bed nets in the future than they do today. And so, on the margin, this is certainly a question that we're incorporating as much as we can into our current assessment of the cost-effectiveness of AMF's activities because we think it's not unlikely that two years from now significant additional funding becomes available to nets and that then would, other things equal, reduce the cost effectiveness of a donation that a donor makes today for nets because it could have otherwise been filled.

0:13:06 EH: So that's the somewhat detailed at AMF update. The bottom line is that we still expect to recommend AMF as our top recommendation to individuals and we're seriously considering whether we should make a big push to encourage donors to give more to AMF now, especially if donors basically know what they plan to give in December and have the flexibility to give sooner rather than later.

0:13:34 EH: There aren't other major updates across our other top charities: GiveDirectly, the Schistosomiasis Control Initiative, and the Deworm the World initiative, which is run by Evidence Action. Though I'm very happy to take questions about any of them. There are updates for all, none as significant as the AMF ones. One note to flag is around some new information that we expect to come out about deworming in a few months. As you may know, one of the key pieces of evidence supporting our recommendation of deworming programs is a randomized controlled trial run in Kenya in 1998, and this is the trial that found impacts on attendance in the relatively near term and then 10 years later found significant impact on earnings for the treatment group. The group that was randomly selected to receive the deworming pills relative to the control group.

0:14:29 EH: The update that's coming is that the researchers have now collected an additional round of data which is the 15-year follow-up to this original study and the researchers have this data in hand and they're in the process of analyzing it and we expect to get an update from them at some point in the next few months. And it's possible that that data could lead us to significantly update our view of the evidence for deworming because of an additional round of data from the single most important piece of evidence for this program that we recommend. So I'm gonna pause here and just open the floor up for questions on top charities and then in about maybe 15 minutes or so, depending on the questions, I'll dive in with another quick update on GiveWell experimental. All the way in the back corner.

0:15:25 Speaker 3: Okay. You mentioned that the efficacy of... The insecticide on the nets reduced the nets' effectiveness by 20% to 30%. Could you talk a little bit about how, on an epidemiology level, that effects how effective a campaign would be for AMF in a particular area. Maybe there are other concerns that [0:15:50] \_\_\_\_\_. So maybe that translate to say they need to reach 5% penetration for a population to reach the same level of reduction [0:16:02] \_\_\_\_\_.

0:16:04 EH: Yeah. So, the question is about the impact of insecticide resistance on our estimate of AMF's effectiveness, and then tell me if this answer is responsive. AMF and its distributions are going for universal coverage, so their aim currently is to cover every single person in an area with bed nets. And so the way that we are thinking about the reduced effectiveness due to insecticide resistance is some reduction in the effectiveness of nets. So if in the absence of insecticide resistance they prevented, they reduced malaria mortality or all-cause mortality by 50%, at a 20% reduction that would be, instead we've modeled it at 40% and that would end up decreasing the cost per impact that they achieved.

0:16:55 Speaker 4: Did I understand you correctly to say that AMF, with like \$40 million, AMF could cover basically... There were two caveats, but they could cover... They could fund distributions almost anywhere, with these two caveats of places that are really hard to reach and places that don't have a lot of need. Does that mean all the kids around the world who are susceptible to malaria are gonna have nets or what exactly is that mean?

0:17:21 EH: Yeah. So it's a good question. I'm just gonna repeat some some of these questions for the recordings. So the question is just when I talked about AMF and its relationship to the global net gap, does that mean that with 40 million more dollars every single person who could be covered by a net who needs a net would be covered. So the caveat that I mentioned is important here and it is not the case that \$40 million covers every single person who's at risk of malaria and doesn't have a net. The biggest exclusion from this calculation is a gap of approximately 100 million nets which remodel as about \$450 million in Nigeria, and we've excluded Nigeria from this calculation because our impression is that it is a hard enough place to work that additional funding targeted at scaling up bed nets in Nigeria is not going to directly and easily lead to additional nets being distributed there. And so we don't think... Another way of putting this is we don't see a potential funding vehicle that we would be confident can take in money and significantly scale up bednet distributions in Nigeria. And that incorporates AMF, we've even talked about this a little bit, and also entities like the

Global Fund where we've had a different type of conversation that included this topic. And so what we're trying to do is include in that estimate only places where we believe it is likely that with more money AMF could move forward with it's model.

0:19:00 EH: I think if someone were to be able to make a proposal that would give... That we would have some confidence in that could scale up bed nets in Nigeria or other places that would be hard to work, we'd be very excited about that giving opportunity, and then similiary... So Nigeria is the biggest example, there's a few others that are particularly hard places to work, and then there's several countries which each have a total funding gap that is much smaller than AMF is currently able to prioritize given their capacity constrains. So I think there's... I'm gonna get the numbers wrong, so Natalie do you know off the top of your head?

0:19:39 Natalie: Right, it's like 5 million to 10 million nets.

0:19:41 EH: So it's like 5 million to 10 million nets across like 10 or so countries.

0:19:47 Natalie: I think it's less than that, but yeah.

0:19:48 EH: So I think basically you have a situation where AMF is not currently prioritizing net gaps that are less than 2 million nets in a country, because of how much funding they're trying to distribute and they would rather spend... They're a small staff, so they're rather spend their staff time on bigger distributions where they can have more impact with the money they have. So another type of giving opportunity that we'd be interested in is someone who could basically do what AMF is doing but for those much smaller distributions 'cause those are ones we're excluding currently from that bednet gap calculation.

0:20:22 Natalie: We've only looked at Africa, I were we think that bed nets are much more cost effective than Asia or Latin America.

0:20:32 EH: Yeah, right here.

0:20:33 S4: I have a question about the [0:20:34] \_\_\_\_ month to month distribution, but is it your impression that the spiky nature of funding where a huge proportion comes in December is harmful to organizations' planning ability and ability to finance things, and would it be preferable in general, not just in this case, for distributions to be right across the year, or is it specifically like we have a higher level of trust in AMF, and we want them to have more money sooner [0:20:58] \_\_\_\_.

0:21:00 EH: Yes, so the question is whether the general spiky nature of donations is harmful to organizations' planning? I think, as a general rule, it's not particularly harmful because most organizations, GiveWell included, to some extent plan based on an expectation of future funding. And so we can go to donors, and especially donors that we have a long standing relationship with and say, "Do you expect to give in December? Do you expect to give the same amount?" And we basically can assume that some portion of them will fall out but if not for that, we can plan based on those gifts. And in AMF's case it's particularly challenging, because they would have to commit so much funding ahead of time relative to what they would ever expect to take in in the absence of a large GiveWell recommendation to our donors which includes Good Ventures, that they really can't just operate based on their best guess about how much money will come through. But for the other organizations on our list, GiveDirectly, SCI, Deworm The World, in talking to them a general impression is that the spikiness of funding is not something that's causing them a major problem. In addition to that, on the margin, I think for every organization, it would be better if their incoming cash were totally spread equally across the year and it creates some weirdness in how on plan, it makes it a little bit more challenging, but my instinct is that it's a particularly bad problem for AMF. Yeah.

0:22:34 Speaker 5: So one question and one follow up to the earlier question about resistance. With respect to the resistance, I didn't quite hear this but, you know, with vaccines you worry about herd immunity, and I've seen some mention of externalities of nets that if fewer people are being infected that has spillover benefits to people around them, and you were saying that you're modeling the difference with this resistance as linear, which would seem to mean that you weren't worrying so much about effects like that. So I ask is that something that you've considered in the models or not? And then second, in relation to the resistance question, do you have reason to think that the areas where new distributions would be happening would have the same levels of resistance because there has been less deployment of bed nets beforehand. And then my question is, so AMF got a very strong recommendation from GiveWell last year and got \$40 million [0:23:47] \_\_\_\_ got the top recommendation relative to GiveDirectly, but it seems GiveDirectly had a far

greater intake of money more than twice as much. I was wondering what does that imply, why are so many of GiveWell's audience going in that direction that is somewhat different from GiveWell's recommendations and do you have an understanding of what's going on there?

0:24:16 EH: Yeah, got it. Okay, two questions, let's just take the insecticide resistant one first. I don't know the details of the adjustment that you're talking about. Josh, do you wanna answer that?

0:24:29 Josh: Yeah.

0:24:30 EH: Josh is the one who is been leading the insecticide resistance work.

0:24:34 Josh: Yeah. The 30% rough reduction in the effectiveness of bed nets is coming from a quantitative model that tries to basically look at correlations between the level of insecticide resistance in area and the number of malaria cases and predicts from there what reduction in malaria cases and then ultimately the malaria deaths will be and so I think that the model should be taking that sort of thing into account, but I'm not totally confident that it's doing it in the idea way.

0:25:07 EH: Yeah, I'd say right now we are mid-stream in terms of updating our cost effectiveness assessment for the insecticide resistance in a more rigorous way than we did it in the past. And most of the way that our final cost effectiveness assessments comes out is all the staff dig in or many staff dig in to many details and so far this is right in the process of being analyzed and thought about and hopefully we'll have a firmer update by the end of the year.

0:25:42 EH: Yeah, I guess the other question Karl you had about resistance was do you expect same or different levels of resistance in new areas versus areas they've gone into the past, and my understanding is like we don't know enough to say and I don't think the areas they are planning to go into are primarily areas that I have never been served by bednets in the past. I don't know why the areas that have gaps today like may have not had gaps in the future but I think the right expectation... And Natalie or Josh, tell me if you think this is wrong... Is that our best guess is the same level as the locations from which we are building our estimate of the reduction and effectiveness.

0:26:29 Josh: Yeah that's right.

0:26:31 EH: Okay. Then the other question was looking at the comparison between AMF and GiveDirectly, we recommended AMF, our number one recommendation to donors was just give to AMF and then donors gave significantly more to GiveDirectly. I think the headline numbers distort things a little bit because Good Ventures gave a significantly... One very large special gift to GiveDirectly earlier in the year but the bottom line is true nonetheless that even though our top recommendation was to AMF, other non-Good Ventures donors I think still gave about the same amount or more to GiveDirectly than they did to AMF last year. I think that I'm not totally certain what this implies I mean one fact is that...

0:27:16 Natalie: Hey Elie, would you just mentioned what GiveDirectly does?

0:27:20 EH: Yeah, sure. GiveDirectly is a group that just gives direct cash transfers to very poor people and so their idea is instead of us thinking we know best what will help people the most in very poor parts of the world instead give people cash to allow them to decide what to spend and the evidence that this program is effective, the people use money productively, they don't use money on things that we all might worry about like alcohol or tobacco consumption. There is very good evidence that they use it productively and very little evidence that they use it on temptation goods like alcohol or tobacco.

0:27:56 EH: I there's different ways to read this, one fact is that two donors collectively gave more than \$10 million to GiveDirectly and that also skews the numbers but I think of the charities we recommend, I think GiveDirectly is attractive to a segment of GiveWell's audience and that segment is... One type of GiveWell fan puts a lot of weight on organizations that are extremely effective and run like excellent organizations and I think as an organization GiveDirectly is the best organization that I feel like I have seen in the 10 years that we have been doing GiveWell and the reason that we recommend AMF above GiveDirectly is because we think the program, bed nets, has significantly more impact per dollar than the program cash transfers.

0:28:46 EH: But I think the tension between giving to an organization that runs a program with high cost effectiveness versus giving to the organization that is run like an excellent organization is one that has been debated significantly at GiveWell over a long time and I find myself very pulled towards the desire to give to the outstanding effective

organization and from talking to donors I think a lot of them feel that pull as well. I think another pull to GiveDirectly is that even though in some ways giving cash is the most traditional least innovative form of charitable giving, it is... The existence of a group like GiveDirectly is it self-innovative because that type didn't exist before and I think GiveDirectly has some of that cache that comes from being a new exciting organization that's trying to do something very new for development in the way that AMF I think to many people, they have heard about malaria, they have been in the news for 10 years, and it doesn't quite seem like they are getting on-board with something new and cool.

0:30:01 EH: The final thing I'll say is I do think it demonstrates this fact that while many GiveWell donors directly follow our recommendation, a significant portion of donors just give money directly to us that they ask us to regrant at our discretion. Many other donors are using our research and then deciding on their own where they like to give instead of just following the recommendation that we make. Yup.

0:30:27 Speaker 7: To what extent do you think... I'm gonna reference that [0:30:30] \_\_\_\_ the notion that you're not deciding on what to do with the money and the people who are in need decide what you do with the money. DO you think that influences people at all towards GiveDirectly? Did you hear that through giving surveys or anything like that?

0:30:50 EH: Yes. So the question is whether people are influenced by the lack of paternalism of giving to GiveDirectly. You're letting the people decide where to give. I definitely think that's a huge influence. I definitely think that that is causing people to want to give to GiveDirectly. I think in some ways, cash is so understandable to donors in a way that... It's like intuitively understandable from our day to day experience in a way that giving a pill that treats a parasitic infection that could cause an increase in earnings 10 years down the road is not really anything that we have formed intuitions about. Cash, just with this easy relationship to our day to day lives, donors have very strong intuitions about it before they ever engage with the evidence is that one of the things that we see a lot of in conversations about GiveDirectly are these two camps forming. I think I called it the anti-paternalism camp that loves GiveDirectly because it's this approach to aid that allows recipients to decide fully what they do with the money that's being given to them.

0:31:55 EH: And then on the other side, there are lots of donors and there was a lot of push back when we first recommend GiveDirectly from people who said, "How could cash possibly be one of the best things you could give to... Clearly, people will spend it on bad things," drawing analogies to poverty in the US or even just assuming that people couldn't possibly make a decision that would be as impactful for themselves as something like a bednet or a deworming pill. I think one of the things that I see a lot more in conversation with donors about cash is they're coming to the conversation with intuitions about the program which are absent in discussions of bed nets or deworming.

0:32:35 EH: So I think now is a good time to pause and I'll share the second half of my update and then happy to take questions on... More questions on top charities, more questions on the GiveWell experimental stuff or any other questions related to GiveWell. So the GiveWell experimental work is work that we started a few years ago and it grew out of this recognition that GiveWell's top charities have not changed fundamentally in the last few years. Our top charities has basically been bed nets, deworming and cash transfers. And the question was, what could we do to create or to find more top charities? We have been, I have been looking for the same type of charity for nine years now and I think that there is some reason to believe that if the organization already existed, we would know about it, we would've talked to them. Over the last few years, a lot of the work that we've done on GiveWell's research is revisiting organizations that Holden or I had talked to back in 2008 or 2009. By and large, those groups have not been more successful than they were back in the early days of convincing us that we should recommend them.

0:33:51 EH: So if you believe that GiveWell has found most of the organizations that we're going to find as potential top charities, the one possibility of a project we can undertake is to try to support the development of future GiveWell top charities, because in order to be a GiveWell top charity, other funders have had to come in at an earlier stage to make it possible. So you need to have a funder who supported the research that created the evident space for an intervention. So we're talking about bed nets, have a strong evidence of effectiveness. These were often philanthropically-funded research that someone paid for that enables us to recommend AMF today. You need someone to give a group funds to get their program off the ground. We came to give directly at a relatively early stage but we didn't give them their start up funding. That's not something GiveWell does because it looks for an organization with a track record. And so some other funder has to be ready to do that. And then finally, we require organizations that have a substantial amount of monitoring and evaluations to demonstrate that their program works.

0:34:57 EH: And so that's also something that you could imagine an organization that runs a great program, and needs more money, and is operating but just doesn't have the capacity to monitor itself in the way that we would find compelling. And so we've made several grants. And when I say we, what I mean is that we've made recommendations

to Good Ventures and Good Ventures have made these grants to some organizations, we think, could potentially be future top charities. I've talked a little bit about this in past events and what I want to do here is just give one update on a grant we made in the past and then just talk about a couple grants that we're considering as part of this GiveWell experimental work. So one of the grants that we're considering is to a group called Results for Development, also known as R4D. And their goal is to try and scale up treatment for pneumonia in the developing world.

0:35:53 EH: Pneumonia is one of the leading killers of children globally and unlike some of the other diseases that we often talk about: Malaria, tuberculosis, HIV, there's no dedicated funding stream for pneumonia. And so, there are cases that this program is underfunded relative to the potential impact. Their goal in this project would be to try to increase coverage for the antibiotics that treat pneumonia, and they're focusing now to start on Tanzania. And they're aiming to do everything from buy antibiotics... I think it's amoxicillin that treats pneumonia because there is not sufficient supply of it in the country... But also improve the distribution from the central points out to clinics, and also improve health clinic practitioner practices so that treatment is distributed appropriately when people come in seeking treatment.

0:37:01 EH: I think the biggest open question here for us is what type of monitoring and evaluation we'll be able to get. This program is significantly less straightforward than the types of programs GiveWell has recommended to date. So, with a deworming program, you basically once a year have to ensure that a child takes a pill, and then the program is over. In a treatment program, you have to make sure that the health clinics, that individuals will come and visit when they're sick, have sufficient supplies in stock, and that the practitioners are correctly diagnosing the conditions for which people come in. And there is some research showing that practitioners can make the wrong diagnosis when they're trying to decide whether a patient has a fever, a plain fever, something like a cold or a flu, malaria or pneumonia.

0:38:00 EH: And so, this is one of the biggest things that we have talked a lot to them about in our conversations on this project. If you're very interested in this and are looking for more information, we have several sets of conversation notes on our website, and as we move forward, we will continue to publish more detailed updates. Our goal with this is that if everything goes in the way that we hope is ideal, that it could eventually become a GiveWell top charity at some point in the future, probably not before giving season of 2018. So this is something of a long-term... There's long-term potential not short-term potential, but the idea would be to support a group like R4D to do additional scale-up of pneumonia treatments globally or potentially, depending on what we end up believing about their track record, their scale-up of this activity, and of other types of treatments in other countries.

0:39:00 EH: And I think our best guess at this point is that the program is about as cost-effective as bed nets. I'm looking around to see if the people that actually did this work most intensely think that I'm getting that right, but I think it's roughly on the order of bed nets. We tried to build in all the skepticism that we normally do to our cost-effectiveness estimates. As a general rule, our cost-effectiveness estimates get worse as we learn more, so we probably won't end up at ten times as cost-effective as cash or roughly as cost effective as bed nets, but that's our current view of the program.

0:39:37 EH: Another grant we're considering is to a group called IDinsight. They're another group I've talked a little bit about these events in the past in the context of GiveWell's experimental work. They're a group that's trying to conduct what they call decision-focused, rigorous evaluations for decision makers. So a lot of the randomized control trial evidence that we rely on and is conducted is highly academic. And that means it's conducted by people who are Economics professors, they have incentives for conducting the research that have in the past led to great outcomes for research but are also not entirely aligned with the interests of GiveWell as a potential funder. So they're very interested in individual academic. It has career incentives that don't necessarily align with GiveWell's interest in assessing a program's effectiveness.

0:40:32 EH: They also tend to be extremely expensive and focused on... Not necessarily focused on the context in which an organization would eventually scale-up the program. So they might test an education program in India, in Ghana, and then you'd have to wonder about the representativeness of those results to the scale-up of an education program in another country in the developing world. So IDinsight is very focused on producing low-cost, randomized control trial evidence that can be useful to decision-makers and organizations about which programs to implement. So the charity implementing the program on the ground, to policy-makers deciding what to fund, and to funders like us who are trying to think about what programs to direct funding to.

0:41:21 EH: And we're considering a grant that would be unrestricted funding to them with the goal of allowing them to significantly scale-up their capacity 'cause our understanding now is that a lot of their work is... They're constrained by their own capacity and are not able to take on as many projects as we would like. And so, they're really, as far as we can tell, a very unique group in this space. There's really no one else that we know who is like them, and so we're seriously



considering this grant in hope that it gives us this additional tool to provide additional monitoring evaluation or randomized control trial evidence about programs we about in the future.

0:42:01 EH: And then the last one I wanted to mention is just a grant update. So, a couple years ago, we made a grant to a program that Evidence Action was trying to get off the ground. Evidence Action is the parent entity of the Deworm the World initiative, and the grant was enabling them to scale up a program that provided incentives to farmers in rural Bangladesh to migrate to urban parts of the country during the famine season and there was a randomized control trial that showed significant effects of this program in Bangladesh on earnings and then on re-migration a year later without additional incentives.

0:42:48 EH: And so we had previously made two relatively small grants to them for some baseline research to determine whether there was a possibility of scale up to consider other countries that they could go into and earlier this year Good Ventures made a larger grant to enable them to start work at significantly larger scale in Bangladesh where they're currently working.

0:43:17 EH: The next famine season as far as I remember is this fall, so the goal is to enable them to run a larger scale version of their program this year, assess that program and then hopefully be in a position to consider them as a top charity next year. I think that still seems relatively unlikely, that all of the data will be in in time for us to make a confident recommendation to individuals at the end of 2017, but we think there's a good chance that we'll be considering them closely in 2018. And our best guess for their cost effectiveness now is... And we often, when we talk about these things in terms of the multiple on cash transfers which we use as something of a benchmark for giving that would be effective and we think we could scale up relatively quickly without significant problem.

0:44:09 EH: So we think this knowing season program is roughly on the order of five times as cost effective as cash transfers with all the caveats of it still being relatively early in the program's life and these cost effectiveness estimates being highly subject to many debatable assumptions. So, that's the update. I wanna pause again and we have about 10 to 15 minutes or so for questions on that or top charities or anything else related to GiveWell. Ah, Robbie.

0:44:46 Robbie: How does the randomized control group dramatically reduce the cost of actually doing a good evaluation and producing those results?

0:44:58 EH: Yeah, so the question is how does IDinsight reduce the costs of randomized control trials. So my understanding is that one of the biggest costs of a randomized control trial run by an academic is surveying enough indicators that we not only have an understanding of whether the program worked in the context in which it was implemented, but we also have some ability to theorize about the mechanism through it worked. And in order to do that, there is a significant amount of additional surveying that's being done. And so I think the idea behind the IDinsight approach is focusing on a much smaller number of indicators to reduce survey costs.

0:45:50 EH: But I don't think I have the best understanding of exactly how they do that and the extent to which they reduce cost let's say just a rough benchmark, there are I think in these will vary the cost of randomized control trials were very heavily depending on the type and the scale, but I think very roughly a cost of around a million dollars for a standard development economics randomized control trial is pretty normal and for IDinsight they're more in the range of a couple of hundred thousand dollars. So they're reducing the costs by 80% or so. But anyhow, I don't think I have the best answer to that question.

[pause]

0:46:32 EH: Yup, yeah right here on the aisle. Did you have a question? Yeah.

0:46:36 Speaker 9: So I have a question about GiveDirectly. It has two motives. It would be a little complicated, sorry. So, there's a question in moral philosophy that comes up about whether you're obligated to give and what drives it is this counterfactual. So if I give \$200 to Oxfam, will this make a difference in the world, like counter factually, I did it and I didn't do it. And I think the idea is that a lot of aid organizations are still vague and the answer is no, it just wouldn't make a difference. And so, I'm thinking about whether I cop up 2,000 and it won't make a difference, I can argue, "Well, I don't even have a reason to give the money an obligation." I'm not saying that's a good argument to run it like voter's paradox of collective action, and all that sort of thing.

0:47:17 S9: So, that's one motivator, and then the other one is when I talk to people about giving, I think two things I've

put a check on giving is uncertainly about effects, you just give the money, you don't know what happens and then cause all diffusion, you just put it into this giant pot of money and who knows what difference it makes. So I'm curious whether GiveDirectly could set something up where they have a fund where all the money that goes into the fund is dispersed within the month to people living in poverty, like 100%. So this either could go to new people or they could just have the people that already have their phones set up to receive the money and so they will just be topped off. So I can see why this would sort of rub effective [0:47:58] \_\_\_\_ the wrong way, right? 'Cause you wanna say, "Let's just do it the most effective way." But I'm speculating that this would be really intuitively appealing to a lot of people. Like, "You put in \$500, in less than a month that money will go to people living in poverty." So the question is, do you think that's a viable thing that they could do and also do you think that it would tend to draw more funding?

0:48:24 EH: Yes. The question is whether GiveDirectly could set-up a mechanism through which people could give directly to very poor people, very quickly, like on the order of less than a month and would that be viable... Would that be possible and would that be, what was the other? Advisable? Or would it draw more people in? Would it be viable, would it be effective?

0:48:50 EH: I think the short answer to those two questions is that I think they could. I think mechanically it's not a very difficult... I wouldn't expect it to be a particularly difficult thing to do, though I don't have the best understanding of all the steps that have to happen before the money can literally go out the door. But my instinct would be that if you were willing to just pass the money along to someone who is already enrolled in their program and you had their number and were not worried about any knock on effects, you could make it happen and then... I think people are extremely... Donors, and certainly less sophisticated donors, are very engaged by the idea of seeing the direct impact of their work and so I think that anything that helps people see their money go to someone more directly would have a significant effect.

0:49:44 EH: I don't think they would have to set up this type of mechanism in order to get that direct impact effect, and I think Kiva is just a great example of a group that presented themselves as if they were doing what you're describing. And years ago, this is what people thought of when they were going, "Kiva, pick a person, the money goes directly to them and then they pay you back." And the way it was actually working behind the scenes was very different. And this has all been hashed out in public and Kiva's much more transparent now than they were about what they were doing, but I think that you can probably get a significant amount of this benefit, if you just wanna focus on the marketing part or how you present what you're doing rather than literally following through on doing this exact thing. 'Cause my guess is that, and this would be a question for GiveDirectly, not for me, that they would not be excited about this idea just because it would create other potential problems for their credibility in the areas they go into if they're just topping people off very, very quickly with the money that's coming through.

0:51:01 Natalie: The money gets distributed relatively quickly, within a year of donating and possibly much sooner, right? On average already?

0:51:14 EH: Yeah, I actually don't know the average lag, but I think around a year is the right thing to have in mind for the lag of money coming in and going out.

0:51:27 S9: And speed wasn't the... I shouldn't have said a month, 'cause it's more of the making a difference that I meant to focus on so if you give \$500, then it doesn't have to be a particular person, but there will be an effect in the world because you gave \$500.

0:51:41 EH: So, I think people... So just taking away the speed consideration, I think that this is something that is easiest to show for GiveDirectly among the groups that we recommend. I think they could do this, I think they may be thinking about how to do this effectively because it really is relative when you compare them to a group like a deworming charity that's going out and giving pills to lots of kids and it's hard to really follow them. GiveDirectly is going to every single individual before they give them money and after they give them money and that gives them this opportunity to get the specific data if they want, so I think they could do this pretty easily. They could do this if they chose to and I think it's something that they're thinking about as they're trying to increase their fundraising capacity. I do think this, even if it's harder to describe, this same reality exists for other top charities and this is the... Meaning when you give you have an impact on the world, though far less easy to show concretely for a group like AMF, where you're contributing an extra couple thousand dollars to a very large potential distribution that could go out and...

0:53:00 EH: So we could talk about how to think about that but the conceptual underpinnings of the framework we use with room for more funding where we're trying to ask, "Will this organization be bottle-necked due to lack of funding?"

And therefore would additional money be likely to make a difference?" When I think about AMF now, they're a group where I believe they are in the very near future going to be bottle-necked due to lack of funding and more additional funding will make a difference. Now, there's some possibility that the amount of money that donors give collectively will not be sufficient for AMF to move forward with an additional distribution, so there's some risk that you give money and it doesn't end up causing something different to happen. But there's also some chance that your \$1,000 is the one that makes the giant distribution go forward and it has an incredible impact per dollar. Yep?

0:53:52 Speaker 10: I was curious to hear more about the question you talked about with major donors closing funding gap versus going part of the way and then waiting to see if other people come in and finish it. Are you guys, like looking at other examples of situations where that has happened? Where a major donor has come in and paid part of the way and then other people are motivated? I guess like, and obviously that question was raised for a reason, so what sort of led you guys to start speculating about that?

0:54:20 EH: Yes. I think the sort of substitutive question that we face is, what should we recommend Good Ventures give to our recommended charities? And Good Ventures plans for giving are such that, given the current funding gaps of our top charities, it would be able to fill the full gaps if it chose to. And so then the question is if, we, GiveWell, consistently recommend that Good Ventures fill these gaps? Does that mean that many of the donors who in the future would otherwise be interested in giving to these types of giving opportunities, will basically never come and learn about us, never be interested in them, because they assume that there's this giant funder that's available to close the gaps. And so today Good Ventures is able to close those gaps. Our goal as an organization, and this is true for GiveWell and also Open Philanthropy, is to find great giving opportunities that far surpass Good Ventures ability to give.

0:55:22 Natalie: And in the short term there's sort of a relatively limited, I think real potential effect on the amount the donors would give. It definitely is... We end up recommending that we leave a really good giving opportunity on the table. And the hope is that in the long term it leads to significantly more overall giving going to really effective places. It's a game theoretic discussion that we've started engaging in. It's not something where we feel like we've reached the final best answer by any means. We've been talking to an economist who focuses on game theory to try to get a better understand how professionals think about these questions rather than us. But one of the things we've found consistently in charitable giving is, we've just devoted ourselves to this question of how can you give most effectively, and we bump into questions that it doesn't seem like people have asked directly before. Ah yep, Mark?

0:56:25 Mark: Can you update us on New Incentives?

0:56:27 EH: Sure. I'll answer this and then one more question, and then we'll break before the open fill part of the program. So New Incentives is a group that provides conditional case transfers in Nigeria. And they're one of the first groups that we funded, as part of the GiveWell experimental program. And the idea is they were trying to get their program off the ground. They were implementing a program, cash transfers, that had strong evidence of effectiveness. And in our interactions with the founder, we felt that, she had a similar commitment to GiveWell, to transparency and monitoring and evaluation. So we felt like this was a pretty good bet for us to take on a potential future GiveWell top charity. They now have been operating for a little bit, year and a half I think, more or less, with the program that they're running. And the program they started with was using cash transfers to incentivize HIV positive pregnant women to come in to clinics and get the AIDS treatment drugs, that would prevent mother to child transmission of HIV.

0:57:42 EH: And in assessing this program's cost effectiveness, we wanted to think both about the cash transfer impact, but also the impact on reduction in cases of HIV. The biggest update on them over the past year, is they found that there was not a sufficient volume of potential recipients of that transfer, HIV positive pregnant women, that they could run a program at a reasonable cost for prevention of mother to child transmission of HIV. And so they added in an additional focus on high risk pregnancies, and they basically said if pregnant women meet one of these criteria, we will consider them for inclusion in our cash transfer program, and the incentive will bring them back for additional prenatal visits, delivery in a facility, and then postnatal visits. And the major outstanding question for us now, is when we look at the combined program, cash transfers some portion of people receiving this, being incentivized towards prevention of mother to child transmission of HIV.

0:58:57 EH: The other focused on at risk pregnancies. What's the overall cost effectiveness of this figure? And how does it compare to our other top charities? And we're in the midst of trying to assess the question of at risk pregnancies, and the impact on maternal and neonatal mortality of women delivering in clinics. And that's not an investigation that we've completed. I think, one obvious question that might come to mind is, "Well, regardless aren't they at least as good at GiveDirectly on their cost effectiveness?" And I think it's true that they, on the sort of straight programmatic level

new incentives, and there's some caveats to this, but they should be as good as Give Directly, because I guess GiveDirectly is cash, plus research. New Incentives is cash, plus HIV plus program, whose effect we have yet to fully analyze.

1:00:00 EH: But in a world where we felt like New Incentives and GiveDirectly were indistinguishable in terms of cost effectiveness, we would end up recommending that donors give more to GiveDirectly rather than the riskier proposition of continuing to scale up New Incentives. One final thought on New Incentives is... Our expectation... Svetha, the founder of New Incentives, she really has met and exceeded my expectations for ease of communication, and ability to answer our questions, transparency, commitment to monitoring and evaluation. They have some of the best programmatic monitoring that we've seen for many charities. So they just have... They use Google spreadsheets to track their progress, and we have access to the Google sheets. So we can just go in anytime we want and see how the program is running, and that's not true of any of our other top charities. And so, in the event that we end up deciding that conditional cash transfers with the conditions that they've set is not something that we're excited to recommend. I can easily imagine us funding Svetha to do something else in the future because she seems like she has now demonstrated the ability to get a program off the ground and run it successfully in a very difficult place to work. And so, that's another possibility that goes beyond just new incentives as an entity. Yup, last question.

1:01:34 Speaker 12: Sort of, as a follow-on to what you were saying at the very end there, are there any other holes that you commented that GiveWell Experimental has identified in the sector where you wish there were organizations for instance, what would you want her to do next?

1:01:47 EH: Yeah. So, what other holes are there? What else could we imagine a new top charity doing that we found? So, we've launched a blog post late last year called 'Charities We'd Like To See' which tried to address this question, though it was very high level and we don't have a great list of specific things we'd like to see an organization undertake. There's an interesting group called 'Charity Entrepreneurship' which grew out of the effective altruism community, and their explicit goal is to create a GiveWell top charity. And so, you can read about their progress, where they're going through, assessments and programs that they might try to implement, and how they're making that determination.

1:02:32 EH: But I think across the board, the types of... The most likely type of programs I can imagine coming up in GiveWell Experimental, are a program that scales up an intervention that was initially measured in one randomized control trial and seems highly cost-effective, but no one has really taken the steps to bring it to scale as an organization that could take in charitable donations. Or, someone who's able to bring one of the banner global health programs that we don't now see opportunities to fund because the gap world wide is closed. 90% of the people who need this have it, but there's still 10% who aren't being reached because of logistical challenges. If someone can come forward and say, "We can reach the last 10% with routine childhood immunizations and here's how we do it." We talked a little bit about bed nets, the same is true for vitamin A supplementation. So they're some of the most effective programs that are going on worldwide, and I don't believe that any of them are literally fully filled or there's no opportunity to cover more people. It's just a more challenging proposition than continuing to throw more money at the problem.

1:03:49 EH: Alright, so, I think we're gonna now break from this take. You have 5 to 10 minutes or so to catch your breath, get some more food, and then we should reconvene back here around 8:15, and Holden will be talking about Open Philanthropy.