GiveWell Research Conference Call, December 6, 2012

0:00:03 Elie Hassenfeld: Well, hello everyone. Thanks all for joining us. We really appreciate your interest and your taking the time to discuss our research. I want to briefly go over our plan for the call. Because of the number of people on the call, GiveWell staff will primarily be the ones talking. If you have any questions, please email them to info@givewell.org, that's I-N-F-O@givewell.org with the subject line, "Conference call question" and we'll read them and answer them on the call. We'll aim to answer all the questions, but we are gonna try to group them topically rather than answer them chronologically. We have noted the questions people have sent in, in advance, of the call and we're going to address those, as well.

0:00:45 EH: At the moment, all your lines are muted but if we do have time at the end of the call and if there is interest, we'll open the call up for more discussion. To let us know that you'd like to speak, you can just hit star 6 on your phone and then follow the instructions. Doing so is going to alert me that I should unmute your line. When I do so, you will be alerted that your line is unmuted, and you can then ask your question and remain unmuted for the discussion.

0:01:09 EH: We're gonna initially break this call into two parts. First, I am going to describe our process for finding charities and then quickly review strengths and weaknesses for our two top-rated organizations. If you are looking for information beyond what I present, there is plenty on our website. Then, after the presentation, we'll answer questions for a while before coming to topic number two, which is our plans for future research. We're tentatively planning for this call to run about an hour, but we're happy to stay on longer if there is interest. We are recording this call so that we can post it on our website. So, feel free to leave at any time, you'll have a chance to finish listening later. Before getting started, I just wanted to introduce the GiveWell staff, who you may hear talking on the line, so you know our voices. I'm Elie Hassenfeld, a GiveWell co-founder.

0:01:55 Holden Karnofsky: I'm Holden Karnofsky, the other GiveWell co-founder.

0:01:59 Alexander Berger: I'm Alexander Berger, a GiveWell research analyst.

0:02:03 EH: Great. First up is going to be the research process that we followed this year. In past years, we had conducted a broad scan of hundreds of charities. We did not repeat this scan in 2012. Instead, we focused our efforts on first, continuing to track and follow up our top-rated charities from last year. Second, we continued to analyze additional organizations that had seemed promising, but were not top-rated at the end of 2011. And then, third, we looked for opportunities to fund additional, proven, cost-effective scalable programs. In particular, this year, we looked into immunization and salt-iodization programs. We have written more about them, but in both cases, they all kind of initially looked as if there were going to be opportunities to fund them but once we looked into the opportunities more closely, it appeared that someone already had or was about to provide funding. We're happy to get into more discussion with that if there is interest.

0:03:02 EH: But going through this whole process left us with three top charities, the Against Malaria Foundation or AMF, and the Schistosomiasis Control Initiative or SCI which were both our two top-rated charities last year. We've also added a new top charity GiveDirectly, which provides direct cash transfers to very poor people in the developing world. Now I am just going to briefly give some basic facts, strengths, and weaknesses about each organization. Our number one ranked organization is AMF which distributes bed nets to prevent malaria. Bed net distribution has the strongest evidence of any charitable program we've seen, outside of perhaps, immunizations. We know, as much as those things are knowable, that bed nets prevent malaria and save lives and at $5

each, they are incredibly cheap.

0:03:47 EH: AMF is an outstanding organization. Its CEO is constantly pushing to run the best distribution she can and is collecting data that could influence the global malaria community to improve. There is a giant global funding debt for nets. Even though they are cheap, there are close to a billion people who need nets and nets wear out after a couple years, so they need to be replaced. All these factors combine to make AMF a really great giving opportunity.

0:04:09 EH: GiveDirectly, our number two charity, runs arguably one of the simplest programs. It just donates about 90 cents on every dollar it receives directly to extremely poor individuals, and it argues that they're the ones most likely to know how best to help themselves. Now, a lot of developing world country governments have run cash distribution programs, but GiveDirectly is the first charitable organization we found that's just kicking in money and giving it out in the developing world. The government-run programs have been successful. They find significant increases in spending on basic needs like food and no evidence of significantly more spending on items like alcohol or cigarettes.

0:04:46 EH: There are some studies that follow up with recipients years later and find significant returns, roughly 20%, on the capital granted. Like AMF, we think the people behind GiveDirectly are sharp and talented, and there is a lot of upsides in their ability in GiveDirectly, in terms of how it can change the world of international aid, to be more driven by what people need rather than what outside charities think are the best ways to improve people's lives. With GiveDirectly, a lot of the judgment among donors may come down to one's intuition about how likely empowering people to directly buy the things they most need is to be among the best charitable interventions.

0:05:24 EH: Our number three ranked organization is SCI which does deworming. This is a health intervention that is very different from bed nets which prevent malaria. In the case of deworming, we noticed, on numerous clinical trials, that the pills SCI deliver kill the parasites that are infecting the individuals. And we also know that distributing these pills is incredibly cheap at around 50 cents a person. But there is a great deal of uncertainty about the difference this makes to someone's life. We know that you won't die if you're infected, so being dewormed doesn't save your life. And even for subtler health effects like nutritional status, there have been a lot of studies and it's still not clear if there are significant effects.

0:06:07 EH: But there a couple of studies that we've read and find credible, that find enormous long-term impacts to deworming. In one, they randomly assigned some children to get dewormed and some not to. They followed up with the dewormed children years later and found that they were earning about 25 percent more than their counterparts who weren't dewormed. Now, there are some tough questions to answer around how credible and representative this study is, but basically, if you believe in those results, deworming is arguably the best charitable intervention. And if you don't, well, you may not be accomplishing very much. Well, SCI has a strong track record of successfully implementing deworming programs. We don't have quite as much confidence in it as we do AMF. And we don't do the same upside opportunities for it to achieve more than its coordination of deworming children. But, we still see that as an outstanding giving opportunity, so this is an organization with a strong track record doing a very cheap health intervention that could have huge and lasting benefits.

0:06:59 EH: So with that, I think it's time to move over to Q&A. So, just as a reminder, if you have

questions for the time being, please just email them to info@givewell.org, I-N-F-O @givewell.org with the subject line, "Conference call question" and then we'll read them and answer them on the call. To get things started, I think we wanted to address some of the questions that we've seen raised about cash transfers and GiveDirectly, in particular. So Holden, did you want to address those questions?

0:07:30 HK: Yeah, I'm happy to. There's been a lot of concerns raised over our recommendation, the cash transfer organization. And I think a lot of it is just, people have definitely very differing intuitions on this particular thing. So, some people think that it's really unlikely that the cash transfers are spent well and others think that people are likely to make the best decisions for themselves. And I think that there is a strong evidence that cash transfers increase consumption, but we do feel that a lot of how one feels about GiveDirectly is going to come down to where one starts off on all these intuitions. So, there are a lot of angles we can cover on this point and we expect a lot of questions on this. But the questions I'm going start with were the ones we were asked to address on this call, which are the questions raised in a post on the Giving What We Can blog.

There were questions that were specifically about the evidence for cash transfers and their cost-effectiveness. So, there are a lot of studies of cash transfers and they generally show encouraging results. They show increased consumption, they especially show increased consumption on food and then there is, basically, a couple of studies that showed more long-term effect, so that people are still benefiting from them, years after getting the cash transfers and kind of implying very impressive rates of return. And so, some of the questions that have been raised are around the quality of this evidence and just as an example, one objection that's been raised, when we say that cash transfers increase consumption, both in the short-term and in the long-term, there's been kind of this observation that, well, they increased consumption according to self-reported data. In other words, in these studies, they're asking people how much they spent and that's the basis on which they're concluding the cash transfers increase consumption and so we've seen concern that these studies are not reliable because of that.

0:09:36 HK: And the main thing I want to observe here is that using self-reported data on consumption or on earnings is really not unusual in academia. It depends a lot on how it's done. And so, if you kind of... If you've given intervention and then you've asked people, "Did this intervention helped you?", that's one kind of self-reported data that we think is very unreliable because the people you are asking, will generally be able to hear what kind of answer you're looking for, and they'll give you the answer they think you want to hear, or at least it's a strong risk. But there's a different kind of self-reported data that I think is much more reliable, which is, for example, you randomize who gets the cash and who doesn't. And then, a long time later, you follow-up with a very long survey that's asking them all sorts of questions about their life and some of those questions are, "How much did you spend last month or last year, over some time period?" And then you go and you compare the control group, the people who didn't get the cash, to the treatment group, the people who did get the cash, and you look for differences in those numbers of the averages people reported.

There are a lot of more obstacles to someone who's trying to tell the surveyor what they want to hear. I mean, they don't know what the control group's going be saying, it's a long survey. It's a very specific question. It doesn't have the most obvious connections to the intervention. And this is generally, how the stronger academic studies are done. Then this is not just about cash transfers. Actually, the best evidence for deworming we're aware of, is using the same basic setup. It's the

kind of the headline figure in our view, the figure that really makes deworming compelling, is this kind of higher reported earnings for people who got dewormed in childhood. But those are reported earnings; that's also done using the same survey data collection process as far as we know.

0:11:29 HK: Likewise, with bed nets, there's this kind of question... So, some of the results with bed nets are not self-reported, right? You can look at kind of clinical records for what happened to malaria. But there is... One of the lingering questions of bed nets is whether people use their bed nets in practice, as opposed to just in the studies. And there again, you've got self-reported data on whether people use bed nets relative to whether they own them. So, there are a couple of things to say here. One is, we do think that self-reported data is not bad, we think it's pretty reliable. It's collected in a pretty good way that's likely to lead to good results. But the other observation is that it is fairly common to collect data this way. And so, if someone is looking at this evidence and saying, "Boy, this is not as iron-clad as I thought it was," well, that pretty much applies to most evidence about the long-term impacts of charitable interventions on people's lives.

0:12:24 HK: For a long time, we've kind of seen a lot of the holes in this evidence and seen it as not so iron-clad and I think that does have something to do with our positive attitude on GiveDirectly and cash transfers. Because in my view, kind of the more skeptical you are of the approach of sort of collecting data in this way, and analyzing it and putting it all together to decide what's best, the more you're going refer to simpler interventions, to interventions of the more intuitive case, like cash transfers. And we don't go all the way. We do think the evidence is meaningful. We don't have cash transfers rated as the best intervention, but it is worth keeping in mind that none of these evidence bases are totally iron-clad. And the difference is, between the cash evidence and deworming evidence, I mean, we think they're pretty much in the same league. So, I didn't address all of the issues that have been raised with the evidence of cash and we're going write more about them, but broadly speaking, we think it's kind of comparable.

0:13:23 HK: Then there's this question of, okay, so let's assume that cash transfers are helping people, but there's this question of how much they're helping people and how does that compare to a bed net or to deworming? And this is another area where, I think, intuitions are going diverge a lot and people are going come to really different conclusions because it is a bit apples and oranges. And it's very hard to find kind of reliable comparisons and it's also very hard to estimate how much good is accomplished by these interventions. And so, with bed nets, you're saving the lives of children under five, you're also reducing cases of malaria. With deworming, you're not saving any lives, but you may be kind of... Or certainly, we don't have strong evidence that you're saving substantial lives, but you may be helping people to earn more later in life, which averting malaria may do, too. And then, with cash, a lot of times you just don't know the exact nature of the benefit that you're giving, and so, it can be hard to estimate it, hard to compare.

0:14:24 HK: We're going be writing about the cost-effectiveness of these three things. And basically, we've tried to compare them a bunch of different ways. And so, we looked at kind of simple measures that leave a lot to the imagination, like the cost per year of being covered by a bed net versus the cost per year of being covered by a metal roof, which is how a lot of people spend the cash. We've also tried to do sort of estimates that take a lot more analysis, and work, and guesstimating on our part and then leave with something a little bit more satisfying, like the cost per life saved equivalent. One of the comparisons we've done for trying to compare cash to other things, is we've looked at... Deworming helps people earn more money later in life. I mean, that, to us, is the most compelling reason to support deworming. Is it better to, instead, just give them the money

and let them invest it, or is it better to pay for the deworming? Which one results in more money later in life, relative to how much money you spent in the first place?

0:15:24 HK: And when we do that analysis, I do think deworming comes out looking better. And we do think that, in a vacuum, deworming is better, is probably accomplishing more good per dollar than cash, and bed nets are probably accomplishing more good per dollar than deworming. But these are pretty uncertain and you can sort of argue it either way. There are reasonable scenarios under which cash is the best. And so, we don't think it's such a blowout that it just resolves the situation. And when we factor in kind of the organizational factors, the fact that we are more confident in GiveDirectly as an organization, the fact that we do see more potential for innovation and for kind of larger-scale impact, that's why we have GiveDirectly as number two and SCI as number three.

0:16:13 HK: And then the final point I want to make, just on this initial discussion of some of the points people have raised about cash, we do not fully embrace the argument that just because cash gives people choice, that's it's better. We see a lot of reasons that that is not enough. So, we've written about these in the past. Alexander actually wrote a post back in May listing quite a few reasons, and not just poor rationality, why it might be better to kind of decide on people's behalf how to help them. One of those cases is sort of externalities. So, bed nets, they don't only protect the person under the bed net, they protect other people because they're killing mosquitoes. And so, it's kind of an externality, it's a market failure argument.

0:16:57 HK: There's also more obvious arguments, like people in the developing world may not have the education or the information to spend the money as well as we could, on their behalf. So, we think there are definitely arguments both ways. But the argument that people can decide for themselves, I think, is an argument, is a powerful argument, that bears consideration. And when we kind of look at all the factors, it looks to us like cash is really one of the more promising interventions out there in the grand scheme of things, one of the more promising interventions that there's room for a donor to expand. We don't think in a vacuum that it's as good as deworming or as bed nets, but when we look at organizational factors, that's why we come out with GiveDirectly at number two.

0:17:42 EH: Alright, so we've had a few questions come in. And so, we're going go through those. We're also happy to return to the discussion of GiveDirectly and cash, or other top charities as we move through the call. So, we had a question come in about Partners In Health, which is an organization that GiveWell had as a number two ranked organization back when we initially started, and one that a lot of donors ask us about. So someone's question was, what's our current take on Partners In Health?

0:18:14 AB: Yeah, I'm going take this one. This is Alexander, again. So, we visited one of Partners In Health's local hospitals in Malawi during our site visit last year. And we've spent a long time talking to our Partners In Health staff members over the years. We continue to have a positive impression on them as an organization. But our big question has been this issue we call "room for more funding" which is, how does a marginal dollar to Partners In Health get spent? A lot of the story that they tell about their impact is that they're really causing changes at the international level through some of their advocacy work and being used as a sort of proof of concept or an example of how really high quality healthcare can be done extraordinarily well-income settings.

0:18:59 AB: That's pretty different from just scaling up bed nets or something that is like very linear, straightforward and simple. It's not the kind of thing that we're super confident in our ability to assess but separately from that, I'd just add that, a lot of the things that they seemed to want to spend more money on aren't necessarily things that are closely-tied in with more advocacy. And so, it seems like the case for their impact has been advocacy-oriented but more money might go into providing more advanced healthcare in the settings where they work and that seems less cost-effective to us than the kinds of health interventions that we're currently recommending.

0:19:46 EH: There's one other thing I would just add about Partners In Health which is that, I think sometimes organizations run programs that are difficult to monitor and certainly, there is types of impacts that Partners In Health could be having that are going to be tough to quantify. Nonetheless, I think it's reasonable to expect organizations to be generating and sharing the data that helps to support a compelling argument that they're having the types of impacts they hoped. And that's something that I feel like we have yet to see from Partners In Health.

0:20:51 EH: Alright, well we had another question come in which is related to SCI and deworming. And this question was whether there is any further evidence that we expect to come forward about the impacts of deworming or SCI that could cause it to rise or fall off your list. So, I'll just quickly jump in and then, maybe my colleagues will want to answer as well. One of the things I think could or would make a big difference to our view of deworming would be additional data from randomized control trials that helps to better ground the effect that is found in the single RCT that we know of. Unfortunately, as far as we know, there is no other RCT going on right now and, were one to be funded today, it would obviously take a great deal of time until the long-term earning data was in. I think given the current state of interest in funding deworming, there is some question about whether that type of research would even be valuable.

0:22:00 EH: But it does add to the kind of the overall picture of deworming that it's the case that a single additional RCT could really make a huge difference if it found no result or it found a significant result in trying to form a reasonable estimate of the likely impact of a deworming program. On the SCI front, as we are recommending additional funding to SCI this year, we're going to continue following them, we hope to continue to gain a better understanding about their activities and how they use funds. We have been following SCI reasonably closely since early to mid-2009 and so, I wouldn't say that I think it's extremely likely that our work on SCI yield a significantly different understanding of its activities than we have in the past but certainly, it's possible and if I had to guess about the most likely place for our view to adjust in the coming year, it would come from improving our understanding of how SCI has and is likely to spend additional funds. Holden or Alexander, anything you wanted to add to that?

0:23:19 HK: No.

0:23:21 EH: Alright. Another question we had coming was about our cost-effectiveness estimates for bed nets distribution. So, the question was, what drove the increase in the cost per life saved for AMF from $1,600 per life, which is what our estimate was earlier this year, to $2,300 per life, which is what we're saying now? And then, also, was there any progress in measuring other benefits of malaria prevention beyond infant mortality reduction? Holden, do you wanna take that one?

0:23:55 HK: Yeah. So, on the first question, why did our estimate of the cost per life saved for AMF go up? So, we wrote a post about this and I think it was called something like "Revisiting the

Case for Bed Nets" and basically, we had a GiveWell staff member go through and do a thorough audit of our research on bed nets and one of the issues that we noticed in the course of doing this is that, the life saved from bed nets are coming not just from malaria deaths, but they're coming from all deaths. And so, the kind of studies in the '90s that found that bed nets saved the lives of children under five, what they actually found is something like a 20% reduction in mortality from all causes in children under five. And we think this makes sense because malaria can weaken you and can make you more susceptible to other diseases, and we discussed this with malaria experts and this checks out. But something that seems worth observing is that child mortality has fallen a lot since the time these studies were done.

0:25:09 HK: And so, we try to make an adjustment for the fact that it's basically, it's a different baseline now. There's fewer children dying in the first place. Now, it's possible that the deaths you could've saved with bed nets have now already been saved by other means, and so, it's actually, at this point just from this factor alone, I'd say it's possible that bed nets aren't saving any lives. It's also possible that now there's children who wouldn't have had a chance before, but now if you protect them from malaria, they'll live, and so it's possible that bed nets are saving more lives. But to us, the kind of middle ground was to look at the fact that childhood mortality had declined and know that you're gonna be working from a lower baseline when you're trying to save lives. And so, when we did our best to adjust for that, that's how we got to the higher number and that is gonna be in our write-up on bed nets and that's also in our blog post.

0:26:00 HK: On the other question, the non-mortality benefits of bed nets. So, there's basically one study that is very similar to one of the studies done on deworming. That is, it is not an experiment. It's a very detailed examination of historical data, and it claims that the benefits of reducing malaria are real and lead to higher incomes down the line. Much is with deworming. It was not really something where you could directly compare the impacts to deworming, because we're talking all these different countries, different settings and it's hard to adjust and figure out how the malaria control they were discussing compares to bed nets. But, broadly, there're a couple reasons to think that bed nets have major benefits other than saving lives. That study is one of them which I don't think would be very strong evidence by itself.

0:26:52 HK: The other case would be the kind of just facts that bed nets are bound to have a lot of short-term health impacts. They have impacts on measures of nutrition, measures of anemia, and furthermore, bed nets are having these effects in children under five which there is other evidence from other domains that health improvements before the age of five quite plausibly pay off later in life and can lead to real differences later in life. And so it's a pretty intuitive case, it's a pretty rough case, and yet I think it's a very strong case. I actually, to me, if I were betting, I would probably be more confident that bed nets have impacts later in life than I would about deworming because of this under five issue that. There's more instances of that making a difference later in life when targeting children under 5. And so, I think there is good reason, I think there's benefits.

0:27:51 HK: As a shorthand, when we've done our cost effectiveness analysis, we assumed those benefits to be similar to those for deworming. That doesn't make a huge difference in the analysis the way that we've done it. If you place a lot more weight on those benefits so those become what's truly important to you, then you have to ask yourself how plausible you find that case and how much you believe that. So, that's where we stand on the non-mortality benefits of bed nets. And, no, there wasn't really any new information on that front this year. I mean, I think something that's important to know is that, if you're looking for long-term impacts on people's lives, these are just

really expensive and really difficult to measure. You have to follow people in the developing world for a long time. It takes a long time to get this information. It takes a lot of money. Any little thing going wrong can make a study a lot less easy to trust. So, if you're looking for long-term impacts, you have to work with the kind of evidence we've been discussing.

0:28:48 EH: And then, there are a couple of other things with bed nets that, given that we're currently seeing AMF and bed nets as our top-ranked organization and the most promising opportunity to give to, we haven't investigated fully. And so there are other diseases are transmitted by mosquitoes. One example is lymphatic filiariasis, another neglected tropical disease. In many places, it's controlled using drugs in a way similar to what SCI does for deworming. In some places, that's not possible because of certain diseases that I believe the treatment that would be given out for lymphatic filiariasis would have adverse effects if it came into contact with this other disease and so, they just control lymphatic filiariasis with nets.

0:29:41 EH: We haven't tried to incorporate that type of estimate into our cost-effectiveness estimate. Similarly, we've written about this very briefly on the blog. We haven't considered the mortality effects of bed nets on adults. And that's something that the most robust long-term data doesn't show that there is a very great problem of adult mortality from malaria in the developing world. Nonetheless, there have been some estimates, in particular, one that came out either late last year or early this year, which argued for a much larger burden of adult mortality to malaria. And if those numbers were true, you'd consider the estimate that we have for a AMF's cost-effectiveness as being significantly underestimating its cost-effectiveness. So, that's about the cost-effectiveness and impact of nets.

0:30:39 EH: We had a question come in that relates to the room for more funding for nets at large. The big picture with nets is that even though on an individual basis, they're very inexpensive at $5 a person, because of the number of people who need nets and the fact that they wear out, the data about how long nets last is not particularly reliable, but the best estimates are that nets last between two and four years, so they need to be replaced pretty often. This means that the global net gap is quite large. I think it's on the order still of half a billion to a billion dollars of funding that is needed just to buy the nets and get them out. One of the things we've done a lot of this year is spoken to many people who play different roles in the malaria control community. And one of the things that came across in those conversations, and these conversations are posted on our website on the conversations page, but it's just the fact that the need for funding to buy commodities is a very real one in malaria control, meaning that there just isn't enough funding to buy nets and more money would buy more nets.

0:31:56 EH: And that, the response that we got from the malaria community was quite different than the response that we got, for example, the immunizations community and this is also something we've written a lot about. That's a case where certainly, the commodity of immunizations are largely paid for and it seems like the other costs associated with delivering the basic immunizations are also there. So, basically, both with AMF, specifically, and with nets, generally, there appears to be a very large funding gap and one that we're not yet close to closing.

[pause]

0:32:39 EH: Another question that I was hoping to address, and this goes back to an organization that we recommended a couple years ago, and that organization is VillageReach which was our top-

rated organization from the middle of 2009 until the end of 2011. And one of the questions was when we recommended VillageReach, there seemed to be a great deal of donations. And then we took them down and then, obviously, the GiveWell source donations have fallen. And what does that mean for the incentives, for an organization to participate in our process? Does the type of analysis we've done for VillageReach create incentives or disincentives for charities engaging with us?

0:33:29 EH: Overall, I think that what VillageReach has done with us is something they deserve a lot of credit for because they've been willing to open themselves up in a way that many other organizations would not have been. That wasn't just in the period during which they were receiving funding or were likely to receive funding but even after they were no longer a top-rated organization, they've remained very transparent, and that is something that we really have appreciated. And I guess we don't yet know the answer but I certainly believe that the donors who follow GiveWell and other donors of similar mindset really value that. So, I don't think it's harmful to VillageReach. Whether it's ultimately beneficial and whether we should provide additional funding for VillageReach, expressly for the purpose of improving incentives for charities to share information, is something that we have thought about and it's not something that we've reached a final conclusion on yet. Holden, was there anything you wanted to add to that one?

0:34:47 HK: No.

0:34:57 EH: So, I think maybe it makes sense to move on to discussing some of our future research. Holden, you can do that and then we'll come back and take questions on that. We're also happy to answer some of the questions that have come in but we haven't got to yet in this part of the call.

0:35:27 HK: Okay, so, to talk about our future research. First off, a lot of our future research will be in the same vein that we've been talking about, so it'll be looking for more opportunities to directly translate dollars spent into lives improved and that includes looking at the best evidence-backed interventions we see, also some of the more promising interventions. They may not be strictly evidenced-backed but they're stuff like clean water and surgeries, things that we may look at to try and find ever more opportunities for this kind of very direct, simple aid. At the same time, we do think it's very important to branch out and to get into areas of giving that are much higher risk and higher reward. This is something we stayed away from to date.

0:36:10 HK: We think it's harder, we think it seems more complex and we think it can be harder to communicate all of the arguments for and against the charity to an individual donor. But we also really we don't want to leave all that opportunity for impact on the table and we want to start exploring other areas. Now, we do expect it to take a long time to reach giving recommendations in new areas. And the reason we think this is because it took us several years to feel confident in our giving recommendations in international aid and looking back, I think it was important that we got a lot of context and a lot of basic understanding, and basic information in this cause. And so, in the future, I think it's very possible that we won't wanna make particular giving recommendations until we've understood the general fields, the general context, pretty well.

0:36:58 HK: Because it's definitely, from experience of ours, that certainly, giving opportunities sounds great on the surface, but the more you investigate it, the more it kind of turns out to be

something very different from what you originally thought. With that in mind, there's a few things we're doing to explore other areas of philanthropy. One is that we're working with charity from Good Ventures to try and get to know sort of the major foundations, Gates Foundation, Hewlett Foundation, understand what they think the best giving opportunities are, why they think that and how they think generally. These are organizations that have been doing what they're doing for a long time and we think there's definitely a lot to learn from them.

0:37:37 HK: Another front that we have done some work on this year and we'll do more work on is what we call "Strategic cause selection". And the idea here is to take broad issues in philanthropy that we think are especially promising. So, it's very common for a foundation to pick issues like US education or environment, and it's also quite common that the way they pick those issues is kind of like following the personal interest and passions of the person funding. And the thing that we wonder is, if you were kind of looking across all the different issues you might choose and looking for the ones where you're gonna have the most impact, could that lead you to a better choice of issue areas than just following your personal interest? And we certainly hope so.

0:38:21 HK: For a couple of things that we've done this year, we've written about what we've read about the history of philanthropy and what's gone well. We've had trouble getting the information about that but we've written about what we've seen. We've also analyzed where today's philanthropic dollars go. And broadly speaking, we look for issues that seem not to get much attention from philanthropy, much dollars from big philanthropists, relative to health-promising matter. We have a list of causes that could fit this criteria and we're gonna be looking into them a little more, and writing about what we find. And so that's kind of having a broad bird's eye view of all the causes we've been trying to take, especially good ones.

0:39:01 HK: And then the third thing that we're doing to explore philanthropy more broadly is picking a couple areas to do more deep investigations of. So, for example, we think it's very important for us to understand the world of funding scientific research, particularly biomedical research, because this is where a lot of the smart philanthropy goes. This is where a lot of philanthropy goes, period. There are certainly arguments that the upside is high enough if you develop a cure for cancer or something like that, that it's worth the giving even at very low odds of success, and we wanna understand more about this field and that's gonna be a major priority for the next year. Another area that we spent a good amount of time on this year and we'll probably spend a good amount of time on next year, is what we call "meta research". And the idea here is that in the process of looking for what works in the international aid, we've reviewed a lot of academic evidence.

0:40:08 HK: Sorry. Elie, Can you hear me?

0:40:11 EH: Yep, go ahead.

0:40:13 HK: Okay, so a lot of the evidence that we review, we've kind of seen ways in which it doesn't seem optimized to help humanity as a whole. And a really simple example is that we'll look at a study and we'll have questions about it, and we won't be able to get the data without emailing the author, which can be a very long process. And this kind of fact that people generally don't share the data and code from the studies and allow others to do other analyses and to check what they did, what's kind of a surprise to us and it's kind of emblematic, that in a lot academic research, you have researchers following incentives that make sense for them but those incentives aren't best

necessarily aligned with what would the best for humanity as a whole.

0:40:57 HK: And so, there is this question, could we as a funder make leveraged, find leveraged opportunities to improve the incentives of academics, to kind of change the way things work, so that there are different norms, different practices and that research as a whole is done in ways that are more conducive to creating value for humanity? And so, in some ways, the work we're doing on these causes, it's a lot like the work that a lot of new philanthropists do when they're kind of trying to get to know the cause. There's a lot of informal conversation, informal exploration, but obviously, a big difference is that we don't see our clients as being kind of one donor. We see our clients as being people like you, as being individual donors, as being anyone who wants to learn from what we're doing. So, we have a commitment to transparency to sharing what we learn as we learn it.

0:41:47 HK: And that's really what our core mission is, is that we are trying to find the best giving opportunities and we're writing about it. And in so doing, we're basically trying to create more of a public conversation about how to give. Because these days there's relatively little discussion of how to give and most of it is just kind of superficial and focuses on why it's important to give generously. And to us, by going through this process of getting to know more and more causes, and writing about what we find, we're hoping to stimulate a higher quality conversation around them and put out a lot more information to help people give well, even outside of the causes that we've kind of originally made our name in.

0:42:34 EH: Alright, well if anyone has questions about the new direction or new causes, specifically, please send them in. We still have a few questions, most of which relate to our current opportunities which we can address. There's also a couple of questions about GiveWell's strategy more broadly, which we'll address. And then again, we're happy to stick around at the end of the call and try to have a more open discussion once the kind of formal Q&A is over. But just to get back to the questions, so a question that came in was about GiveWell's current thoughts about providing restricted funding. So, restricted funding is where one goes to an organization and says, "Please, you can have this money but you can only spend it on a certain activity," whether it's a particular project or a particular type of program like distributing bed nets or deworming, or something else.

0:43:38 EH: Historically, GiveWell has always recommended donors provide unrestricted funding but earlier this year, we wrote that there maybe times where GiveWell would recommend providing restricted funding. So, I think there's two major issues to be thoughtful on that with restricted funding. The first one is truly mechanical, wherein a smaller donor goes to a larger organization and restricts the funding, there's just a question of fungibility. This could happen to a larger donor as well, but I think larger donors have more ability to direct the project that's going to occur. Let's just go pick on a large charity, UNICEF, and says, "Here's a $1,000 donation, I want you to spend this on malaria," there's certainly the possibility that that doesn't cause $1,000 more to be spent on malaria activities by UNICEF. It can spend unrestricted funding that it had allocated to malaria on something else and then allocate this $1,000 to malaria and because of fungibility, the amount ultimately spent on malaria doesn't change.

0:44:50 EH: There's this practical question. There's also a kind of a conceptual question about whether donors should have their own priorities and believe that they know better what an organization should be doing than the organization itself. And I think this is a somewhat difficult

question to answer. And certainly, the less familiar one is with an area, the more one should be modest about when things will happen and how confident one is in their own views. But I think there are certainly cases where restricted funding makes sense. One case in particular, where we've thought a lot about restricted funding is a case where the organization is so large that it's almost as if it encompasses many smaller organizations.

0:45:41 EH: So for example, UNICEF, it's such a gigantic organization, that is almost as if thinking about restricting funds to malaria versus nutrition is, in a sense, like slightly less restrictive than it might seem, because UNICEF itself may not have a particular set of high level priorities and in some ways, may even prefer donors to go to the area that they are most interested in and fund that area. There's also ways to potentially get around this problem of whether one is trying to drive priorities or be receptive to an organization's priorities, something that we often try to do in our initial discussions with organizations that run multiple programs, or even let's say, within UNICEF's nutrition program where they have multiple nutrition projects, is first ask, "What are your priorities? What are the things that you most want to fund?" before we come in and set down our own ideas about what should be funded.

0:46:49 EH: If one is able to answer that question of fungibility which I think is hard for small donors, easier for larger donors and then, can be careful about this question: "Do you have sufficient reason to believe that you got a better sense of the right priorities than the organization?" I think there are cases where restricted funding makes sense and it's certainly something that we do intend to continue looking into, especially as we try to find more of the intervention that we characterized as proven, cost-effective and scalable, meaning those programs that have really strong evidence of effectiveness and appear to be very cost-effective, and may have the ability to incorporate more funding to scale them up. Many of those may require providing restricted funding. For example, providing funding to UNICEF for some sort of health program. In those cases, it's certainly something that we're going to be considering because it seems like an opportunity to do a lot of good and learn about how to accomplish good in aid.

0:47:56 HK: The short answer to the question is that we're looking. Right now, we're not putting that restriction on ourselves. So, we're looking for ways to deliver the most evidence-backed interventions. And whatever that looks like, we're looking for it. And so, when we say we're having trouble finding opportunities to fund immunizations, that includes restricted funding. And so, basically we're open to it and it hasn't turned into giving opportunities yet.

0:48:29 EH: So, we have a question about the funding for nets, which is, "If it is a billion dollars or less, why has it not been filled already, given that it is a small net amount relative to all aid and health spending? Things like this is low-hanging fruit that should have picked."

0:48:47 HK: Sure, so, I'll take this one. I think one thing that's important to note is that there are other proven interventions, and when we say nets are the most proven intervention we found, that does have a funding gap. And so, there's a lot of money going to immunizations. There's a lot of money going to a lot of things that are proven. And so, that's where some of that aid goes. As for the rest of it, why the gap? I mean, malaria funding has gone up a lot since 2004, it's gone up really rapidly, but it still hasn't reached the point where we can get universal bed net coverage across the board. And I think the question why that is would be kind of a question about international politics and about how the aid agencies work. I don't think that the allocations are always following the same criteria that GiveWell follows.

[pause]

0:49:44 EH: Okay. We also had a question about...

0:49:48 HK: Sorry, can I add something on that last question which just occurred to me. So, I mentioned internationally, then I mentioned kind of the aid bureaucracies and the aid agencies, and those are to me the natural funders of something like bed nets because it's delivery of something proven. And a lot of people might be wondering why the big foundations don't do it, like why wouldn't the Gates Foundation come and close this gap? I think it's important to note that private foundations often believe that they can accomplish a lot more good with high-risk projects than by filling these sorts of gaps. I think that's a very common belief among private foundations. It might even be universal. And I'm not sure whether we agree with it but I think that's part of the reasoning behind this one and will expand into other causes, that it's something that there's investigation and it's certainly seems to be a common belief that these had a very direct proven interventions, are not actually the best way for start money to get the most impact per dollar.

0:50:52 EH: We had a couple of questions about a couple other areas of international help aid. One question's about what's the current status of our looking into anti-retroviral therapy, which is the drug, the treatment for HIV. The reason that we had deprioritized that in the past was because of cost-effectiveness issues, that the cost of those drugs were quite high in annual terms. And what they did was treat people, not cure them, and for that reason, we deprioritized. AIDS, in general, it's an area that it's certainly on our priority list, I'm not exactly sure when we'll get to it, but we've had some initial conversations with some people recently about it as we tried to get around it and think about the right way to look into it further. So, I'd say at the highest level, it's an area that we looked into briefly a couple of years ago but we definitely hoped to come back to it because it seems like it might be an area with more of the proven cost-effective, scalable global health programs.

0:52:02 EH: And then, we also had a question about tuberculosis. The question was, it seems like there are many ordinary TB cases. So by ordinary, I think the question means regular tuberculosis as opposed to drug-resistant tuberculosis. Drug-resistant tuberculosis is harder to treat, it's much more expensive. But the question is that "it seems like there are many ordinary cases that could be treated with the normal accepted WHO recommended program and these cases are going undiagnosed at the moment. Do you have any interests or leads in that area?" So, tuberculosis is an area that, again, we've looked into over the years and we did continue to look into this year. The complication with tuberculosis is that it's not clear what the right path is to improve tuberculosis care. It doesn't seem like there are gaps in literally providing the drug that people take for their tuberculosis treatment.

0:53:10 EH: So, one of the approaches that organizations take, and an organization we looked at this year that was along this line, there's an organization called "TB REACH". TB REACH is a grant-maker that seeks to find smaller organizations that are trying to, essentially, encourage people who may have tuberculosis to get tested, so that they enter into the standard country tuberculosis system which includes both the public and the private sector. And when we looked at TB REACH, there's two big questions that we have about this activity, neither of which we were able to answer in a way that gave us confidence. One question is, "Do the activities that these organizations undertake increase the number of people who enter the normal system of tuberculosis care?"

0:54:04 EH: And just to explain what I mean by the normal system on tuberculosis care, in many places, tuberculosis drugs are available in the private sector. And so, when people become sick with tuberculosis, it's not that they just die because they don't have drugs, instead they go to a private pharmacist who may not properly diagnose them. They probably, though we don't really know, don't take the full course of tuberculosis treatment, but because one's tuberculosis symptoms tend to subside after a relatively short period of taking the drug, the people, at least in the short term, probably stay alive and their tuberculosis doesn't kill them.

0:54:45 EH: And so, there is this question of, well, then, what is the humanitarian impact. Even assuming the organizations are able to bring individuals from what's called the unregistered tuberculosis care system into the registered tuberculosis care system, what is the humanitarian impact of that act? It, at least, appears to us that there is limited evidence that it has an effect on directly saving people's lives. It may have an impact on reducing the spread of drug-resistant tuberculosis because one of the ways in which drug-resistant tuberculosis spreads and expands is through individuals who are infected with regular tuberculosis not taking the full treatment regimen. Now this is something, again, that we have looked into, we've tried to understand the frequency with which a treatment default, meaning someone who doesn't complete their regimen, turns into a case of drug-resistant tuberculosis.

0:55:49 EH: This is something that has proven somewhat difficult to model. There are notes on our website where we've gone through the conversations we've had where we've tried to, but the bottom line is we haven't yet been able to arrive at an estimate, in an understanding of this cause, such that we feel good about recommending an organization in it. Again, this is an area I could easily imagine us coming back to, especially by talking to either more people who are well-versed in the field, such that they could convince us that the types of programs these organizations are running, are having the impact that they aim to. But that is not something that we've done yet and not something that's high in our priority list for next year.

[pause]

0:56:50 EH: Alright. Another question we had was about this philosophical question that we've alluded to in some of our blog post and writing, which is the value of interventions that improve lives versus interventions that save lives. Holden, do you want to address that?

0:57:14 HK: Yeah. I mean, I think this is something where there is definitely not internal agreement on GiveWell and our views shift a lot, even person to person. So, personally, I value improving lives much more then saving lives at least for the lives of children under five, and there's definitely widely diverging opinions on this. One observation that I want to make is that a lot of these philosophical debates aren't necessarily that essential for resolving a question of where you should donate. And so, to give an example, I mean, if you want to improve lives, I think the case that any of our three top charities improves lives, they each could be defended as the strongest case.

0:57:54 HK: So malaria, I think you've got these short-term health impacts that are tangible and there's very suggestive evidence that, like deworming, it might have these longer-term impacts and I find that pretty creditable, and personally, that's the one that I would bet on as being the most convincing case. For deworming, the case is there for improving lives and not for saving lives, and in cash there's also a case that you're empowering people to get their own returns, make their own decisions. If you really value saving the lives of five-year-olds a lot more than you value improving

lives, and that's like really the main outcome you're going for, then I think there is basically no case for any charity that we've written about, other than AMF. I mean, the case that bed nets saved lives is very strong, and anything else we've seen with a strong case for actually saving lives, we haven't been able to tie it to an organization as a funding opportunity.

0:58:51 HK: And we've written about some of this. I think there's a post by Natalie Stone a while ago on this topic. This is a different philosophical debate: reducing suffering versus increasing empowerment. She's more interested in reducing suffering, I'm more interested in increasing empowerment. But a lot of these debates are not that essential to answering the question of where you should donate, and so it's a matter of what are the organizations we're looking at that are carrying out highly promising interventions in ways that are highly transparent to us and in ways that we're confident in.

0:59:26 EH: So, we have one more question that I think is on the topic of the current top charities and other associated interventions. Obviously, we're happy to keep discussing them if more questions come in, and then there are some other questions that are not right on this topic. But this question was about the intervention of cash. And what our view is on the possible situation that giving the reason that people don't buy bed nets or buy deworming pills is just lack of information, and if they had the information, people would make different decisions if they had the understanding of what the impacts would be. Now, this, I think, is reasonable. I think it's certainly plausible that people don't have or are not fully weighing all the benefits of health interventions.

1:00:15 EH: Frankly, one of the things that I'm most interested in about, when I think about GiveDirectly, is the possibilities it offers for future research into these exact questions. So, one of the things that I'd be really interested in seeing GiveDirectly do is, at some point, some form of its process, not just provide cash, but provide cash plus information. When people are enrolled in the program, you can tell them about bed nets and deworming pills, in addition to telling them about the cost to replace the roof, or the benefits to migration. There's all sorts of information that could be shared with recipients of cash, and that offers some insight into what's actually going through someone's head when they decide how to spend their money. And so, when we talk about GiveDirectly's upside, this is one of the possible upsides that I see, that there's a possibility that we'll learn more about these types of questions.

1:01:22 EH: So now, I guess I just wanted to get on to a couple questions that were not directly related to our top charities, but other questions related to GiveWell. So, the question on GiveWell, "The organization has incentives that will inevitably sometimes be different than its mission. There are malicious ways for GiveWell to be selfish, and try to grow itself and then, there are more unconscious biases. How much thinking does GiveWell do about these types of biases, and how to avoid being influenced by them?" I'll just answer this quickly, and Holden, you can add if you want. This is something that we're very aware of, and we work hard to try to be aware of the ways in which we might be biased in the things that we're doing.

1:02:11 EH: I think that the research we do is straightforward, and is not biased. Nonetheless, I think the most important way in which we deal with it, is just by being totally transparent about what we do and why we do it. For our top charities, what we've looked at and what evidence and argument had led us to the conclusion that we have, and then, we have conversations like this where people can ask us questions and we put this all out in public. And so, well, obviously anyone can say our work is not biased, we just try to put all the information we can out there, transparently, and

allow everyone to make their own decision about the credibility of our work.

1:02:50 HK: I think it's a real challenge to think about how to make this work of deciding where to give. How to make it more accountable, more reliable, how to take it something beyond what it is today, where it really seems kind of willy-nilly, help people decide where to give. Our selling point is that we will be able to make an argument about what we believe, and that the information we give you will be the best information we have available. So, that's our brand and we expect to lose customers and we expect things to go badly if we fail to deliver on that. And I think that's something that is unique about GiveWell. There are other organizations, like foundations deciding how to give, but they're not out there making a commitment that they will explain to the world at large why they're doing what they're doing, and that people who don't think it adds up should sort of stop being their customers.

1:03:49 HK: So, that's something that I think is different about GiveWell. It's not perfect, it's not as good as like perfectly measuring our outcomes and holding us accountable for those, but that just isn't really realistic. That's where I think we're coming from and what we're trying to bring about. And I think that also relates to some other questions people ask about GiveWell. We have done things in the past where we've kind of chosen to go the route of putting out information that is not necessarily the most favorable to our mission in a short-term sense, and that includes writing about the struggles with VillageReach that includes admitting a lot of limitations of the evidence around bed nets and deworming, and cash transfers. But that's kind of us eating that short-term cost because we're making a long-term bet on this brand of having our arguments be the best, most accurate things we can say in support of what we believe.

1:04:43 EH: Someone wanted to know about and understand more about the relationship between Good Ventures and GiveWell, and GiveWell's upcoming move to the Bay Area. Very brief, and again, we've written a lot about this, Good Ventures is a major donor and I think their giving will grow over time. It's a group whose values about transparency and about cause, agnosticism, which is something we've talked about in the past, are values that they share with us. And so, in a lot of ways, Good Ventures is somewhat like the perfect partner donor for GiveWell. And so, we do work very closely with them in our research, when our interests overlap, which they often do, we work together. When they don't overlap, we go our separate ways. And I think that both Good Ventures and GiveWell are providing a lot of value to each other. GiveWell to Good Ventures, in that we're providing a lot of resources to do research that Good Ventures would otherwise not have. And then, Good Ventures to GiveWell, not only by directly funding our top charities, but also by creating the incentives for some organizations to speak with us who otherwise might not.

1:06:13 HK: So as far as the move to Bay Area goes, that's a combination of wanting to work more closely with Good Ventures and also wanting to test the idea that there's going to be a more receptive audience to GiveWell there than in New York. And I think we have some reasons to believe that'll be the case, it's certainly far from guaranteed but certainly, it feels like our network there was not as good as our network here right now, even though we've been here for five years. And so I think that's a risk worth taking for GiveWell that the costs are actually not exorbitant. We can move back if it doesn't work out but that seems like something that has very high potential benefits. So, that is part of the reason, and part of the reason is to be closer to Good Ventures so that we can work more closely with them. At this time, they have one full-time staffer but we have enough in common in terms of the things we're trying to do that it's helpful to coordinate pretty closely and that is only obviously, in the areas where our missions overlap, that's only when we're

trying to do the same thing.

1:07:14 EH: So, we got a question about some of our future causes, so, Holden already did, but then, why don't you take the answer. So, the question was, "Any hints about what the promising neglected causes are? Any thoughts about causes that focus on the far future? How are we going to deal with regression to the mean in the future given the weaker evidence-base for those projects?"

1:07:47 HK: Sure. Yeah, just a couple of examples of what causes that we have been looking into because we think they're promising for one reason or another. Certainly, the category of Global Catastrophic Risks and something that a lot of people have raised to us is something they think is promising and we think there is kind of a sound argument here which is that Global Catastrophic Risks are things they could have just like such a huge cause with humanity. An example would be climate change. Another example would be risk of being hit by an asteroid. It's hard to articulate who else besides a philanthropist has really got the incentives lined up to work on these things. And so, it's kind of a natural fit for philanthropy, it's the kind of thing where even a small chance of having an impact could be enough to kind of make it worth it because the potential impact is so big.

1:08:42 HK: And maybe, there maybe areas that are relatively neglected within this but the last part is something now we're not sure of at all. And so, certainly, the dollars spent on these causes from our rough cut looks smaller than the dollars spent on a lot of other causes, but we're not actually sure whether when we go looking out there for the biggest threats, large-scale threats whether we are really gonna see a lot of low-hanging fruit or whether we are going to see that there are appropriate agencies and people working on these things. And so, that's a big question. Kind of a first global catastrophic risk we've been looking at and this is a little different from the other ones, is climate change. That's something we'll be writing about fairly soon. And that one is not one of the more neglected ones in philanthropy terms but, certainly, is one of the ones that has more evidence out there regarding how serious it is and so, it is an interesting place for us to dive in and start looking around.

1:09:38 HK: Another cause that's pretty different that's an example of the kind of thing we are looking at, is migration. There was a study that was kind of arguing that an intervention to help people migrate and give them the liquidity and the risk protection they needed to migrate from rural to urban areas, was very beneficial for them and actually led to them migrating again the next year without getting that same kind of support. So, this is an example of something that we think has a strong intuitive case as one study but this isn't something where we think it's time with that intervention to be looking for a charity that's rolling out some specific migration intervention, rather the thing we're wondering is whether this is an anti-poverty intervention that just isn't getting any attention. And maybe there should be more study of it, maybe there should be more discussion of it leading down the road to a stronger view on how promising it is. So, that's a very different endeavor from looking for a charity that will translate dollars directly into vaccines delivered or something.

1:10:39 HK: I do believe that when your evidence is weaker, that is a major issue and I think that, all else equal, something that has strong evidence behind it, is a better bet than something that's got a lot of guess work. And so, I don't have much to say about how we're going to handle that except that we're just going to make our arguments about why the promising-ness of a certain area is enough that it makes sense for us to be recommending things there, despite not having kind of proof that what worked before will work again. And I do think we'll be able to do so because I think that one thing we believe is that, we believe in investigating things very thoroughly and investigating

them from a lot of angles.

1:11:23 HK: And so, I think any recommendations we make in these broader causes, there will be a lot of conversations, a lot of conversation notes with people who know what they're talking about, and we'll be looking at a lot of different angles and hopefully, if we recommend something, we'll be able to make an argument that, "Hey, this thing is worth trying. It makes a lot of sense and there is no one else doing it and we got a good organization to do it." And so, our recommendation can be made despite the lack of evidence that it worked before and it can work again. But time will tell, we don't know yet and we'll see as time goes on, how strong a case we're able to make.

1:11:59 HK: One thing that's worth noting is that all the proven interventions that we talked about today, those wouldn't exist without a lot of high-risk efforts that people have taken. So, we've got a lot of medical technologies that were developed through research that was, at the time, there was nothing proven about it. And so, that's one thing that makes me think there's probably more to good giving opportunities than just taking the things that people have already taken the risks to establish, and just paying to them to reach the maximum people possible. It seems like there are pretty good opportunities on the other end, as well. But that's not something I know, that's just something that I think is worth checking out.

1:12:37 EH: Yeah, another cause that someone asked about is about political advocacy. You could imagine something like improving or increased spending on aid, or other type of political lobbying. Again, that's the type of activity that certainly seems like it could be among the most impactful charitable activities. It is one that many philanthropists focus on. It's also challenging because the issues themselves are complicated, and it's not always clear what side to be on. It's complicated to measure impact, so it's one that we hope to look into, but that is in the somewhat far future of GiveWell.

1:13:19 EH: So, we've now covered the questions that have come in. If anyone has questions that they just wanted to discuss directly, we're happy to stick on the line for a while and do that. So, in order to do that, you just have to hit star 6 on your phone then, follow the instructions. I think, you then have to hit "1" and then you'll be in the queue and you'll have a chance to discuss. So, we'll give people a couple of minutes to see if anyone wants to do that. As always, feel free to drop off the line when you want, if you had enough, we are planning to post the recording of this call on our website so you'll have the chance to hear it later. Great.

1:13:58 S?: And my question is, is it seems like you've done a lot of work thinking about if you were wrong about these charities being the most effective, what would that look like? You did a full audit of AMF and I'm sure you do a lot of thinking also in your spare time, what's the scenario that makes our analysis wrong? I'm curious five years down the road, if you could give me a guess of how likely is it, five years down the road you'll be like, "Yep, we nailed that one, that was definitely the right decision to make given what know at the time." And how likely it is that it would be, "Well, it was a pretty good recommendation, it was a very good gamble, we did some good, but there are probably other ways we could have gone there." I'm not sure if that question makes sense.

1:14:48 HK: Yeah, I think that's an excellent question. I think certainly we do spend a lot of time thinking about what it would look like if we were wrong, and we still see ourselves as new beginning. I don't look back and think that the charities we recommended early on didn't necessarily work out great. Again, I think they were kind of good bets, as you say. So, as far as where we stand

on today's top charities, I think on this pure question of how likely are we to look back and say, "This roughly did as much good as we were thinking it would do," I think there is a pretty steep difference, and I'm going to speak for myself, let others chime in if they disagree, I think a pretty steep difference between number one and between the others. I think AMF is doing one of the most proven interventions, and we've seen it carry out this intervention, and we have a rough sense of what's going on, even there, there are definitely concerns.

1:15:42 HK: I mean, there's questions about whether AMF is going to find partners that can execute the way that AMF needs to execute, and in 2012, it didn't find those partners at the rate that we would have wanted and didn't find enough to get the money out the door. So, there's that kind of question. There are these questions about bed nets that are lingering. There are these questions of whether people use them. I think the answer is "yes". There are questions about insecticide resistance, I think the answer is that bed nets are still working. But certainly I think these could turn out to be wrong. And so, I think that I don't know. I think it would take me more thought to put out a percentage I feel good about, and I'll probably put in that thought. But I would say that AMF, I think, there's a pretty high chance that we'll look back and feel like it did what we hoped it would do but a high chance, not overwhelming. I think with the other two, there's a lot more question marks, and I think they're kind of calculated risks.

1:16:37 HK: So, with GiveDirectly, it's a newer organization and part of the reason that we're recommending it is this kind of a calculated risk, it's that we see the upside. And so, I think there's actually a very substantial chance that in the next few years, they'll discover that there's something unexpected about the way they're doing their program that calls for changes, and so the money spent on transfers in the first couple years was not what we hoped it would be. In addition to the issues we've raised with deworming, I think I wouldn't be surprised if a replication of the study didn't find the same results. I think there's also this issue of the organization where we haven't always felt totally clear on what the money's doing and on exactly what's going on. And so, I think there's a risk that we look back and say, "Well, this money wasn't spent in the way that we kind of imagined it being spent and the activities carried out were not exactly what we imagined being carried out."

1:17:38 HK: So, the short answer is, I think, there are decent chances with all our top charities, and this is not something where I'm saying, "Yeah, we got this." And that's why we value learning so much, and that's why a major weight in how we're ranking these charities and the fact that we're recommending them, is about, if we move money to these charities we think we will learn more, and we think we will get better at giving in the future, and we do think that most of the impact of our recommended giving, the direct tactical impact could be in the future but the impact of giving now is that we learn and then we get better at it. So, that's a major factor on how we're thinking about all this stuff, it's a major factor at how we rank stuff and that is what it is. I think that AMF is definitely something we feel more confident about than the other two, in terms of looking back and thinking that it is what we thought it was. Elie, you're thoughts on this?

1:18:31 EH: Yeah. I mean, that's basically how I would describe it, too. I think if you're looking for the organization with the combination of the highest likelihood of having impact, especially in the context of everything we know, I think AMF is a pretty good bet, given the way they have been behind the intervention and the organization's track record, and the other two have some more of the complications that Holden mentioned. I still think they are great organizations working on great intervention, certainly relative to the all of the different charitable programs out there.

1:19:15 S?: Okay. Thank you.

1:19:21 EH: Great. Here's a question which is quite, that someone emailed in which is quite difficult to answer, which is, "What is the ratio of your perceived impact for AMF versus Partners In Health?"

1:19:39 HK: The ratio of our perceived impact for AMF versus Partners in Health. This is one of these times when I think that putting a number on it, especially an off-the-cuff number, is likely to create more confusion than information. I think it's easier to give a qualitative answer that basically, we just have a lot more questions about what the dollars are doing with Partners In Health, in addition to major concerns about whether the good accomplished per dollar is comparable to AMF, and so, I feel strongly that AMF is a better bet for individual donor-giving than Partners In Health.

1:20:17 HK: I feel strongly that bed nets are probably accomplishing more good per dollar, given the information we have and I feel strongly that I would feel much more confident and needing to give them my dollars for doing something along the line with what I imagined them to be doing, is giving to AMF. But to come up with a ratio, I think given the either one does some good, and I think that given that either one very likely does not do exactly the kind of good that we calculated and imagined. So, I think that's probably where I'd left that one, unless either one of you wants to try and give some more context there.

1:20:50 EH: The idea of providing a dollar per impact type of quantification of what charity to do, is certainly something that we would like to do. It's something that we, I would say, intended to do, as we began GiveWell. And it's just something where over the years, as we've actually tried to deal with the figures, we've just struggled to come up with a way to do that in a way that actually seems to clarify things more than it obscures them. And so, this question of, well, "What ratio would you use for AMF versus PIH?", or similarly, any of our top charities versus other organizations that we submitted, whether it's something like Doctors Without Borders, which we have a generally positive view on, or many other organizations, it just becomes so difficult.

1:21:55 EH: And the way that we've handled that is to say there are a set of organizations that everything we know about them, from the strength of the evidence to their need for additional dollars, to the strength of the organization which includes their commitment to transparency and sharing information, as well as other things that we've written about, that then leads us to organizations that we're going to really strongly recommend, give our money to, recommend donors to give money to. And then, for the rest, we just don't really feel like we're able to give a number that ends up being helpful.

1:22:33 HK: Yeah. So, one kind of attention point on cost-effectiveness analysis in general is that, I do want to note that I feel like we've said a lot of things on this call about how the limitations of cost-effectiveness analysis, and we'd shied away from giving numbers on the basis that that may create more confusion than clarification. I want to put in context that we do think that we have possibly put more effort into getting the right cost-effectiveness figures, like the cost per life saved, than anyone else in the world. We've seen the kind of efforts put together by Disease Control Priorities report and we've seen the efforts put together by others, and I think there's an argument that we have worked harder at getting this done than anyone else.

1:23:18 HK: And so, it's something that we really want to do and it is something that we definitely, consensually, cost-effectiveness is what we want. We want the maximum good accomplished per dollar, and in terms of how we spent our time, we try to get it right and it's not just because of the numbers, it's also because trying to get it right helps clarify our thought process. It helps make sure we're asking all the right questions. And even after all the struggles, I think this is very worthwhile. I think that cost-effectiveness analysis has actually been a really good use of time. It's helped us think about the issues better, it's helped us make sure that we've thought about the issues we need to be thinking about. And so, in the scheme of things, I think we're very pro cost-effectiveness analysis but it is through that experience, and it is through being up close with it, and doing so much with it that we have recognized its limitations, and have come to feel that a lot of times, putting a number out there can create more problems than it solves.

1:24:14 EH: Okay. Well, it seems as if we've covered all the questions. We really appreciate everyone who's joining this call tonight and asking us questions. One of the things we worked very hard to do is to share our research and make it public and engagement from people like you who want to spend an evening with us on the phone, is something that we really value and appreciate. We really appreciate getting questions from our audience and talking about them over the phone or answering them over email or in blog comments or wherever. So please do not hesitate to send those in. We would really appreciate it. Well, thank you all for joining us tonight and have a really great night.

1:25:13 HK: Thanks!