

GiveWell Research Event, San Francisco, December 3, 2012

**GiveWell Research Event, San Francisco, December 3, 2012**

**0:00:00 Elie Hassenfeld:** Well, thanks to everyone for coming. I'm Elie, I'm one of the co-founders of GiveWell and this is Holden, one of the co-founders of GiveWell. Really appreciate you all coming out here. One quick note before we get started. Consistent with our general approach to what we do of trying to be transparent, we're gonna be recording this and we're gonna post audio and a transcript on our website. If there's anything that you say that you don't want to go on the website or in the recording, just let us know and we'll edit it out before we go live. We're happy to do that. We just wanted to give you a heads up that that was happening.

**0:00:52 EH:** The basic program is we're just gonna... We want this to be an informal discussion about the research that we do. So, I'm gonna talk for a few minutes about some of our past work. Holden's gonna also say a few words. We're gonna pause then for questions. We really encourage you to ask whatever is on your mind about the work that we do. And then, Holden will also talk a little bit about some of the future directions for our research. And then, we're happy to stick around afterwards if anyone has any more questions about GiveWell and what we do.

**0:01:25 EH:** So, GiveWell started about six years ago when Holden and I were both working in the private sector trying to figure out where to give our own charity. And when we did that we found very limited information that was helpful. A ton of financial information like how much does the CEO make and how much does this organization spend on overhead but nothing that was answering questions of how well programs really work. And what I mean by that is, let's say you're giving to a program that provides clean water in Africa. Well, you wanna know does anyone go back a year later and check whether the wells that were dug are still working or whether they've fallen into disrepair? Is the water that's in the wells, is it really clean because that's one of the goals of digging these wells?

**0:02:07 EH:** For us, we're really interested in getting the most bang for our buck and so there's other questions you want to ask. How does getting the clean water compare to supporting a malaria program or an AIDS program? And how do those programs internationally compare to supporting programs closer to home like education or homelessness? So, nothing back when GiveWell got started really addressed those questions at all and nothing else today helps to address those questions but that's what we're trying to do. What we do is in depth research on charities to find outstanding giving opportunities for donors. And every year we'll publish a list of two or three top recommended organizations that we strongly recommend donors give to because they're really outstanding.

**0:02:51 EH:** And so, one of the things that I wanna' share is some of the main lessons that we've learned in the years that we've done this. And probably the single most salient thing that we've learned is that very often, charities tell these really exciting compelling stories of the impact that they're having but once you look deeper, the story doesn't really match the reality. And that's not to say that the charities are running scams because most are not. Most of the organizations that we've had any dealings with are well intentioned. It's just that charities are trying to solve some of the world's hardest problems and often they don't succeed.

**0:03:29 EH:** So one example where this came through for us was in the cause of micro-finance. The cause of micro-finances is making small loans to extremely poor individuals and the idea is that they can invest in their business and pull themselves out of poverty. And this, for a long time, was a very exciting part of aid. But when we looked at it, the story is a lot less clear that this is what's actually happening, The type of debt people are taking on starts to look a lot more like credit card

debt. The interest rates that people pay are extremely high. They can be 30% a year, 100% a year, and people often spend the loans on normal consumption.

**0:04:12 EH:** There's been a couple recent really high quality studies of the impacts of micro-finance and there's just no evidence that this has a significant positive effect on people's lives. And so, that leaves us with the question, when you give to micro-finance are you doing as much good as you can? Are you helping people more than you're harming them? Because there's this very real question of harm. Another example, and this is a more positive way in which it can be difficult to give effectively, is causes that are so well proven that large funders like the foundations like the Gates foundation or governments have already stepped in and provided funding to the extent that it's really needed.

**0:04:54 EH:** And this is something that we've seen to some extent in the cause of immunizations. Immunizations are cheap, they save lives. But even though many people are still not reached with immunizations internationally, those who can be easily reached appear to already have been reached by the programs. And so, even though charities are still out there and they'll say, "If you give us money, we'll do more immunization cause that's what we need money for," we haven't found the organization, though we've looked really hard where we're confident that more money means more people reached with immunizations.

**0:05:29 EH:** And so GiveWell's basic value proposition is that we are a full-time staff and we're putting thousands of hours into trying to answer these questions and find outstanding opportunities to recommend to donors. We ourselves have a group of supporters who fund our operations and they both want to promote the cause of giving more effectively but also want to try and change the rules of charity, to change it from a world where charities are rewarded for telling outstanding stories to a world where charities are awarded for accomplishing an outstanding amount of good.

**0:06:05 EH:** So with that background, I want to talk a little bit more about the details of our research. And so I want to talk about the process we use, and the questions that we ask, and then I'll turn it over to Holden to talk about the recommendations for our top charities right now and some of their strengths and weaknesses. So the way that... Aside from a very brief period early on at GiveWell where we were looking at domestic US causes, we've largely focused our time over the last five years on international giving. And that, itself, was a reasoned decision that I'm happy to talk more about if anyone's interested, but the way that we work is we'll cast the net as wide as we can and consider lots of organizations.

**0:06:48 EH:** So at this point we've considered really all of the big name charities that are working internationally, your UNICEF, Doctors Without Borders, Charity: Water and hundreds of other smaller organizations out there, but we're not trying to put a rating on every single organization. We're trying to find those outstanding ones that we recommend to donors. So we'll use flags to highlight organizations that seem like they're worth following up on in more on depth. So one example of a flag is if an organization runs a program that has a really strong evidence that it accomplishes a lot of good, something like immunizations, we'll flag them and then we'll look at them deeper. We have a whole host of other flags.

**0:07:28 EH:** Once we've decided to go deeper, we'll talk to them, we'll ask for documents. And normally at that point we'll say, "We've learned enough now. This organization doesn't seem like it's going to match up with our top charities," and we'll post a page on our website explaining what

we've learned and why this organization doesn't match up. With the ones that we think are most promising, we do a ton of due diligence. We'll talk to staff, we'll review all of their internal documents that get at how they know whether the program's effective. We'll review all of the academic evidence to say what does this program accomplish in general, and then we'll also go and do visits in the field to see their operations on the ground.

**0:08:08 EH:** And so through that process we arrive at our top recommendations. The top three that we just recommended now, there's the Against Malaria Foundation which provides funding to distribute bed nets that protect against the mosquitoes that transmit malaria in the developing world. There's Give Directly, which is an organization that does something very simple. It just provides direct cash transfers to extremely poor individuals with the idea that people who receive the cash may know more about how to help themselves than we might. And then finally, our third ranked organization is one called SCI, which stands for the Schistosomiasis Control Initiative. Schistosomiasis is a parasite infection that affects hundreds of millions of people across the developing world.

**0:08:56 EH:** Holden will get into more specifics, but broadly the types of questions that we're trying to ask about these organizations are things like, how impactful is it to run this type of program? With something like parasite infections, how bad is it to have one of these infections? What does it mean for your life and how you'll develop or how you'll feel better if we're able to treat these infections? How much does it cost to treat them? Do you need to treat them every year or can you treat them once and they go away? These are the types of questions we wanna know about the program.

**0:09:26 EH:** Then there's questions about the organization itself. If you're an organization that distributes bed nets, do you have a process for making sure that you only give nets to the people who need them? Do you check back after to make sure that people who get the nets, actually use the nets? What are your thoughts and are you sensitive to the fact that there may be resistance building up to the insecticide on the net such that they're no longer as effective as they used to be? So these are the questions that we want to know from the organization.

**0:09:52 EH:** Then a final big area that we're concerned with is this question of how the next donation will be used. So you could have this organization that has been really effective but has a new donor, the question is what more are you going to accomplish with this new donation? So for example, you could have an organization that does surgeries, right. So you need both money and surgeons to perform these surgeries. Well, if you have plenty of surgeons and not enough money, well, then additional donations are probably going to accomplish a lot of good.

**0:10:20 EH:** But on the other side, if you don't have enough surgeons more money may not be easily translatable into additional surgeries. And this is something we've seen happen with one of the more well-known surgery organizations. As it took in more money, it started spending on things that were more tangentially related to its mission; not direct surgeries, but things like raising awareness and education. And so with that frame of our process and questions, I guess I'll turn over to Holden to talk a little bit more about our top charities.

**0:10:49 Holden Karnofsky:** Sure. Did you mention cost effectiveness? Okay. Alright, we missed that. So our top three charities, these are kind of what we... We put all our effort throughout the year into finding the charities that basically we would give to, and that we will give to, our own personal

money, and that's kind of the frame in which we're trying to recommend these. As people, who are still relatively new to giving, although now we've been at this for several years, but also serving people who are new to giving and who don't necessarily have all of the context to do soft evaluation, we do tend toward interventions that seem to us to be easier to document, that seem to have stronger evidence bases and that generally seem to have kind of a simpler theory of change in a sense, or a simpler model of how they work.

**0:11:39 HK:** And so these are the charities that we currently recommend giving to. We hope people will consider giving to them, and ask us whatever questions they want about them. So the number one ranked charity at this time is the Against Malaria Foundation. They were also our top ranked last year. And they are distributing insecticide treated nets to fight malaria. Now, insecticide-treated nets are the best proven intervention that we know of within the set of charitable programs that sort of have the kind of funding gaps we've been discussing. So, in other words, Elie mentioned that immunizations are cheap and save lives, and this is true, but we have had a very hard time finding kind of this opportunity to put in more money in vaccinating more children.

**0:12:24 HK:** That's not the case with bed nets. So, out of all the health interventions or all the interventions, period, where we found that there's this gap that you could help fill and reach more people with more dollars, the strongest evidence case belongs to bed nets. There have been 22 randomized controlled trials of them, that's where they kind of give out bed nets to some people and not to others randomly, and then they follow them and they see things like the rates of malaria, and malaria deaths, and measures of nutrition, and measures of anemia, and very consistent results that these bed nets have significant impact on malaria. They save lives.

**0:13:01 HK:** Our current estimate of cost per life saved of a child under five is around \$2,300, which is, if that is right, that's a very low price to save a life. Of course, these estimates definitely have their issues, and we're happy to talk about that. So, that's the intervention. There are a lot of organizations that are distributing bed nets. So, what sets AMF apart? And the answer is basically their transparency. And so, we've talked to a lot of organizations distributing bed nets, but with most of them, it's hard to find out much about the bed net distributions, about the details of when they happen, how they happen, whether people can continue using the nets, et cetera.

**0:13:42 HK:** And AMF is really outstanding in terms of, with every distribution, you are able to track, on its website, when the nets are shipped, when they get delivered, when they're handed out, and AMF is collecting data on usage, data on malaria cases before and after, publishing them publicly. And overall AMF has just been a model of transparency, accountability. They have a strong advisory board. I mean, it's actually one guy within the advisory board who is doing the due diligence, working with partners on the ground to get the nets distributed, and just really into transparency, accountability, and answers our questions really well.

**0:14:21 HK:** So, that's number one. Number two is Give Directly. And this organization is something that we haven't recommended before, it's a new organization, and they're the ones that do direct cash transfers. So, their goal is that for every dollar you give them, 90 cents ends up in the hands of extremely poor people in the developing world with no strings attached. The people can do whatever they want with the money. This is a controversial intervention. A lot of people are instinctively very uncomfortable with it. We've definitely heard from some of our audience on this point, and we'd be happy to hear from you on it too. And people have very different intuitions on it. There are people on the other side -- and actually, we would count ourselves in this group -- who

say that this is an intuitively very appealing intervention.

**0:15:09 HK:** What you're doing is you're basically taking those dollars that you have and you're putting them in the hands of the people who need the help and you're letting them decide how to allocate them. Now, obviously, there are reasons that we would be able to spend money better than the aid recipients. There's public health externalities, where controlling malaria could help more than just the person using the net. There are simple differences in information and education, but there's also a lot of reasons to think that the people receiving the cash may know things that we don't. And so, to us, this is an intuitively appealing intervention, it also has been very extensively studied. Actually, of all the interventions that are not helped, this is the one that we've seen with largest number of high quality studies on.

**0:15:54 HK:** People have asked these tough questions. Are people just spending the money on alcohol? Are people getting incentives not to work? How are people spending the money? And overall, the concerns have not been borne out, although things are still inconclusive, and people tend to spend the money. They certainly invest some of it. A lot of times they may have some surprising opportunities to turn money into more money just because of the fact that they're stretched so thin that having a little bit of free space allows them to do things that we might not ordinarily think of. They spend a lot of the money on food, improving their diet, improving the quality of their food sources. And so, this is what the evidence says, and we basically find this intervention appealing, but clearly there is a discussion to be had, and we wouldn't say the case is 100% clear cut.

**0:16:41 HK:** Give Directly has been around for a year. We think they also are a really outstanding organization in terms of transparency and accountability. So, they have been a complete open book. They're actually running two high quality studies on their program where they randomly choose who gets the cash transfers and follow them. They've published their analysis plan in advance, which is a very unusual thing to do for a study, and it makes it... When the results come out, it makes it much easier to tell, "Okay, is this what you were looking for, or did you go back after the results and kind of cherry-picked the ones you wanted?" So, this is an unusual thing that they've done and we think it's very cool. And they've been around for years, so we've been able to look at how their program has gone so far. We do believe they're getting cash into the hands of extremely poor people. We visited them, as we visited all three of these charities for several days in Africa.

**0:17:35 HK:** And then, our number three charity is the Schistosomiasis Control Initiative. So, they do deworming, and that is treating people with pills that kill intestinal parasites. Now, deworming is an interesting intervention because the pills work, they do kill the parasites. The big question is, what is the impact of the parasites? A lot of the infections are asymptomatic. So, you have people who don't even know they're infected. You wouldn't be able to know they're infected. Some infections are very severe, but often the people with severe infections would go into the clinic and get treated.

**0:18:08 HK:** So what deworming is it's actually treating the entire population, whether or not they have these infections, to get rid of some of the more subtle milder ones. And the question is what is the impact of that on quality of life? And I think the answer is somewhat unclear. So the short-term benefits, a lot of the clinical health benefits are not very well-established. People have done studies on them but haven't found very much. There are a couple of studies that show very exciting long-term effects. So, in other words, 10 years after people get dewormed as children, they're making

substantially more money, like 25%.

**0:18:45 HK:** Now, if you find the studies valid, then deworming looks like an incredible deal because it's 50 cents to treat a child, and that's including all costs, and the benefits are significant 10 years later. It's somewhat hard to compete with that. On the other hand, there are issues, there's only a couple of the studies that we've seen on this question that we consider are kind of worth looking at, and they both have their issues. And so, it's kind of a murky question. And then, the organization has a track record. They have gone into countries repeatedly, they've gotten deworming programs going. They have documented that they've reduced the level of infections and the commonness of infections, and they've been very transparent, and they've shared all the information that we've asked for.

**0:19:29 HK:** We do have more confidence in our number one and number two charities as organizations. We've had an easier time understanding what they're doing. We've had an easier time following them over the last year, keeping up with their progress, and having a sense of what's going to happen next. But they all three are outstanding organization. All three are doing interventions that have a lot of appeal for us, and all three are places that I would give my money. I will be giving more to number one than number two, and more to number two than number three. So that's our top charities. And with that, we're gonna take questions for a while, and then we're gonna talk about our future plans to kind of extend GiveWell and go to areas that are beyond this framework of kind of these proven linearly deliverable interventions. Yeah?

**0:20:15 S?:** How do we build expertise across the huge number of areas that charities cover?

**0:20:21 HK:** Right. So the question is... We're gonna repeat questions because we're recording, "How do we build expertise across the huge number of areas that charities cover?" One answer is that we're not experts in any of these areas, and we're trying to leverage the work of other experts. So, we don't do academic research, we review research. We look at what's already out there and kind of the more we can take advantage of, the better. So if there's a group that does a good summary research as long as we're able to spot-check it, gain confidence in the group and its methods, we'll use it.

**0:20:57 HK:** So really what GiveWell is about and what our research process is about is it's not becoming an expert in every field because that's completely unrealistic, but it's kind of using our time efficiently and it's in a way is what... The way we think of it is that as a donor, it's what a donor would do if they were trying to figure out where to give their own money, if they just had a lot more time and resources than they have. So that's one of our biggest challenges, and a lot of what we do is take shortcuts. So when we started off, we sent out a lot of grant applications. We offered grants, and we figured the charities would come to us with their case. That's saved us some time.

**0:21:35 HK:** And then, as we were going through that type of process, we noticed other patterns that allowed us to save more time in the future. At this time, one of the patterns that we're taking advantage of is, A, that we become fairly convinced that you get the best thing for your buck when you help the poorest people in the world instead of the poorest people in the US. So we focus on international aid. And B, we believe that the interventions that have strong evidence behind them, that we basically identified most of them, and that there's a small number that you can identify them by searching public academic literature, and so we can look for groups that are working on the most promising interventions. That saves a lot of our time as well. Yeah?

[background conversation]

**0:22:34 EH:** Okay. So the question is about... So the British government's Aid Agency is the Department For International Development, DFID, and I guess they're thinking of an approach where they provide aid agencies with funding based on the number of DALYs that they can prevent. And a DALY stands for Disability-Adjusted Life Year. It's a health metric that tries to put saving lives and improving lives in the same term. So if you save a life, you save one full life worth of DALYs because otherwise the person would have been totally dead. If you prevent the flu, let's say, you've saved a very small amount of DALYs because the person would have been sick but not really dead for a short period of time.

**0:23:24 EH:** So for a lot of the people that... In some ways it's conceptually the thing that is very important. You want to spend as much... You wanna accomplish the most, and in some ways that is a combination of maybe improving lives or saving lives for the dollars that you spend. In practice, we found the DALY metric itself, in attempts to quantify the costs and the benefits of programs, as something that has been more problematic than it's been helpful. There are some published DALY estimates where people will say, "Well, if you get to this program it costs X dollars to avert a single DALY."

**0:24:06 EH:** And we've looked at these numbers, we've vetted them and we found really gigantic errors in the calculations. An error without a hundred X in one of these estimates that, the one that we looked at which was the one that was the most cost effective. This is the one we've written all the details about on our website. So, the big picture is that while conceptually this approach of trying to accomplish as much in these terms as possible with the money you give is very appealing, in practice we don't see that the tools are there to measure these things in a way that allows you to actually accomplish what you set out to.

**0:24:44 HK:** So, as another illustration of this we try to do estimates of things like the cost per life saved or the cost per life saved equivalent, and you can see them on our website. We publish the spreadsheets, and the spreadsheets show like with all the different reasonable associations you could make, what's the cost per life saved. And as an example, I mean when we tried to do deworming, we were asking questions like how do you value these improved earnings compared to saving a life? And how skeptical should you be of these studies? And how should you adjust for this way in which the studies were unrepresentative? And in the end we thought, you could make it an argument for a set of assumptions that would point to like \$50 per life saved equivalent, or you could make an argument for something that showed \$20,000 per life saved equivalent, and pretty much anything in between and actually things outside that range, as well.

**0:25:33 HK:** So, we... Again, conceptually, definitely if we could kind of perfectly measure and say this is how much good we're accomplishing, this is how much it cost and you divide them into each other, if you could capture everything that way, then I would give my money based on that because that makes sense to me, but I don't think that that's realistic. And I think things get even more complicated when you use the DALY metric and that's why the problems we're talking about were there. And so, in the end we think that giving is going to involve a lot of judgment calls, and it's going to involve some intuition. And it matters a lot how strong the evidence is, too, right? I don't think that you can just kind of guess every number and multiply them together and go with that. So, conceptually, yes, but we don't really see it happening.



**0:26:21 S?:** Is there generally a strong evidence based for most interventions or is that usually the thing that prevents you from recommending something just because you don't know whether it does or whether it doesn't?

**0:26:33 HK:** So the question is the lack of a strong evidence base, is that the most common thing that prevents us from recommending a charity? And I would say a qualified yes. There's a little more to it than that because we want evidence for a reason and so definitely different interventions have different standards to clear in my opinion. Some charities' programs are kind of very complicated and elaborate, and you want a lot of evidence. For me, anyway, I think that certain things like corrective surgery or cash transfers potentially face a lower burden of proof. But definitely the combination of like, it's not obvious to me that this works, which when we're talking about the developing world, I think it's fair to say that nothing's obvious and a combination of that with a lack of strong evidence, yeah, for most interventions, I think would, we would not be confident in for that reason.

**0:27:24 S?:** So, I mean, could you put like sort of the top two or three things on the research agenda which would help you sort of gain confidence in things that might be interesting?

**0:27:34 HK:** The top two or three things on our research agenda which would help us gain confidence in things that might be interesting. So you mean the things we're going to do to gain more confidence or the things that could exist?

**0:27:44 S?:** Well, if you're not... If you guys aren't doing the research yourselves.

**0:27:51 HK:** Right. Right, we're not.

**0:27:52 S?:** Well, and so the research that other people are doing...

**0:27:55 HK:** Right. So, we need to see more high quality studies, and that's kind of a broad term. One of the kinds of studies that we usually not always find more convincing is the randomized trial, I've described a couple of those, where some people get the intervention, some people don't, and then you follow them both and you can be reasonably confident that whatever you saw is because of the intervention. So we would like... If we saw more of those, and especially something that's very rare, is to see like many high quality studies on the same basic program that outside of health is pretty rare. There's a lot of interventions that have like, you know, one person study them one time and maybe two times, but to see a whole bunch is pretty rare.

**0:28:42 HK:** And then there's some other subtler factors. I mentioned publishing your analysis plan in advanced. I mean, we'll look at a study and we'll say you know, where they found an effect in women between the ages of 21-23, I mean, is that what they were looking for? Or is that kind of them trying to find something? So people publishing their plans in advanced, sharing their data and code because when we really put weight on the study, we really try to go back to the data and pick it apart and see what we can see and we've done that sometimes, so little or things like that to make the studies easier for us to assess.

**0:29:21 S?:** So, Dambisa Moyo once said that, she's skeptical of bed net donations because they put the domestic bed net industry out of business. Is this true and if so what kind of effect does that

have?

**0:29:38 EH:** Yeah. I mean, so the question is whether the donating bed nets for free puts the domestic industry out of business. I think the... To me, you always have to ask this question about which side of the argument there's the burden of proof. And, you know, I think that in the case of giving out bed nets, it's like... This thing, if we get it to people, we're saving their lives and, yeah, I think there is some possibility that by giving things out for free, you are preventing this industry from growing that could on its own be self-sustaining. I don't think that there's, literally, it's happening that bed net manufacturers exist and are being put out of business by the aid industry. I think as you balance this question you want to ask yourself, given that it's a fact that giving out these nets saves lives, which side of the question would require more evidence for me to change my behavior.

**0:30:43 HK:** There is... One way to ask the question is would you dismantle the kind of international bed net distributions for the sake of hopefully getting the domestic industries back up and running and the costs seem pretty huge and pretty tangible especially because of these arguments that bed nets should be given out for free because there's public health issues, because my bed net protects not just me but other people in the village. And so, the cost can be pretty tangible and the benefit, I think pretty speculative and also not as... Well anyway... So that's the way we put it but obviously, people can make their own judgments.

**0:31:20 S?:** And there also were pre-existing businesses like in the case of bed nets, they were massively scaled up because...

**0:31:27 HK:** Right.

**0:31:28 S?:** Of international aid so...

**0:31:32 HK:** Oh, yeah.

[background conversation]

**0:31:43 HK:** Yeah. So the qualitative... The guess for how cost-effective GiveDirectly is. This is something, I think the cost per equivalent life saved is something we have not tried to do. The best way we've come up with for comparing cost-effectiveness for GiveDirectly is to look at the rate of return on cash and compare it to the implied rate of return on something like deworming. The idea is that if you deworm a child then they earn substantially more money and we have a study of saying how much, although it's only one or two studies, 10 years later. And so you can take that money and discount it back and say, "How does this compared to the invested returns they could have gotten if you have just gave them money in the first place?"

**0:32:30 HK:** And there's definitely... I mean, there's studies... The deworming returns are extraordinarily high if you take as face value. I think there are good reasons not to take them at face value. There are also studies showing some really, really surprisingly high returns per cash. So, in the range of like 40%, 60%... Our estimate, 60% a year of invested returns and this may sound crazy, I think it's probably unrealistic but I think probably less crazy than it sounds because I think we're used to thinking of invested returns as something that you have to be very clever and beat the market to get. But when you had no cash at all, it could be something like, "Well, you were renting

a motorcycle by the day, now you bought it."

**0:33:16 HK:** And so that's... You're earning a rate of return on that that's not, like some sort of brilliant investment that out smarts everyone else, that's just a matter of not having any liquidity and then getting some. And so, there are certainly some pretty strong claims for both. I would say that the... If you take things at face value, deworming is much better even if you tried to make like what we think are somewhat conservative adjustments, deworming still looks better and I would probably bet on deworming having a higher rate of return but I don't think it's super clear cut. You can certainly argue and certainly be skeptical enough of these deworming studies that you think cash is better.

**0:33:53 HK:** So, the way that I'd summarize of this I would... Well, another... Okay, so another framework, we've come up with every angle we can for thinking about what's more cost-effective than what? So, another way we think about it is, bed net costs about \$5, deworming costs about 50 cents, is the whole \$1,000 cash transfer per family being spent that well? And intuitively, we would guess that it's not. So, these are arguments that cash transfers are doing less good per dollar than health interventions. That said, I mean, there are other factors in play. GiveDirectly is a new organization, they are a great organization, they've got an evolving model.

**0:34:30 HK:** I think that they are asking interesting questions that are gonna lead to interesting discussions in the international aid world and even kind of challenges for the way aid is currently done, and perhaps replaces an even less cost-effective program with these cash transfers. So, they've got a lot of upside, there's a lot of opportunities to learn from watching them and the more skeptical you become about all these academic evidence and cost benefit analysis, at least for me, the more attractive these cash transfers start to look. So, our basic attitude is, we've ranked at number two, I will give much more money to AMF than GiveDirectly but it's something that we think is not an obvious call.

**0:35:09 EH:** Yeah. At a basic level, there's the question of the cost-effectiveness or the impact per dollar of the program and then at the same time, there's this question about the organization and it's important to not lose sight of the value of an organization that is implementing a program, particularly effectively. It certainly seems that way to us because it's so easy for organizations to run programs but then not run them in a way that we'd see really great outcomes. And so, with GiveDirectly, one of the things, aside from the fact that we think there is a pretty good argument that cash itself is among the better things that you can do, the organization is strong enough that it's one of the better ones that we've seen in terms of carefully running a program and learning from it. And that learning is so rare but yet so valuable in aid that is something that we're really excited about.

**0:36:07 HK:** I want to cut in with why we recommend three charities at all and why we kind of have this recommended allocation that isn't just give all your money to number one. Because I think there is an argument for just giving all your money to your best bet, to the best charity. It's not personal investing, it's not something where you have to diversify and kind of guard your risks. At least in our opinion, it's not that when I give, I want to do the most good possible in kind of an expected value sense.

**0:36:35 HK:** But, the reason that we recommend three and that we hope we'll see a substantial money go to all three is that when you think of GiveWell as a whole, we tracked \$5 million in 2011

that was attributable to our research and hopefully more in the future, more in 2012 which is hard to say because most of the money comes in in December. And when you think about that pool of money, would you rather give it all to one or some to one of each? The argument for getting some of that money to different organizations is that you get to learn more and your giving improves more in the future.

**0:37:09 HK:** So that's a lot of our thinking here and actually I would guess that dollar-per-dollar kind of abstractly isolated deworming does more good than cash. But I think our learning opportunities are going to be better with GiveDirectly. We feel more confident in the organization and we feel that more public goods and more knowledge will be generated by it, and that's a big part of what we're thinking. Because we are still new to this and in some ways we think that some... Perhaps most of the good that I do with my giving will be far in the future when we just know a lot more because of having given early and having learned things from it. Yeah?

**0:37:44 S?:** Could I ask a couple of questions about your malaria program? Which I think is in Malawi, is that correct?

**0:37:51 HK:** So The Against Malaria foundation, currently its major distribution is in Malawi. It's also looking for partners in other countries.

**0:37:58 S?:** Well, couple questions about that, which may have implications to other recommended programs. The first one is about sustainability. The Against Malaria Foundation is distributing these bed nets and carefully monitoring who's using them and what the short term impact is. In three, four, five years, all those bed nets will need to be replaced and then again, and then again and then again.

**0:38:23 HK:** Yup.

**0:38:23 S?:** So my first question is around what do we know and what statements can you make about ensuring that this rolls on the decades or as long as it's needed. My second question is about partial impact. If those communities in Malawi are sleeping under bed nets, malaria will be reduced but it won't go away. People will still get malaria. People will still die of malaria.

**0:38:47 HK:** Right.

**0:38:48 S?:** For malaria to get to zero or zero deaths, as you know, a lot of other things have to be done in addition to the bed nets. So, in your evaluation of that program, are you looking at who's doing all those other things against malaria, and how is that coordinated with the bed net component?

**0:39:09 HK:** Right. So the question is kind of a broad... How much are we looking at Against Malaria Foundation's sort of long term localized impact? So, in other words, they're distributing bed nets but are they working to make sure those bed nets will continue to be distributed in the future? Because they do have to be replaced every few years. And what is their role in the larger effort to kind of eventually eliminate malaria? So this is something that we... I don't think there's a good argument that AMF is kind of, is especially outstanding in terms of it's efforts on these fronts. I think what AMF is doing, and the way I would describe it, is that they're doing an intervention that is having a lot of positive impact, that is long term impact for the people it helps. If a life is saved

from malaria, if you have a childhood that is malaria free, that is a long term impact for that person.

**0:40:05 HK:** They also, by hoping to get the bed nets funded and distributed, they're contributing to the larger malaria effort. Because, now the programs that are trying worry about covering the country with bed nets can think about doing other things instead. Now, we have noted a gap for bed nets so we don't think it's actually fully fungible but within a country, it may be. We think it's important and valuable that these bed nets are a part of the larger effort. But, I wouldn't make any claims that AMF is leading the charge on these other things. And I think that the overall health system has a lot of people doing a lot of important work. It's not all about AMF. What we're trying to do is find a part of the system, that as individual donors who are not closely involved, who don't have the ability to kind of follow these things in softer ways, we're trying to find a part that we can help with. It's kind of a comparative advantage idea. It's not...

**0:41:03 S?:** Could I ask you a very small supplementary question?

**0:41:04 HK:** Sure.

**0:41:06 S?:** In the communities in Malawi where these bed nets are being distributed, do you or do they have any data on the other important elements of any malaria control program, which are diagnosis and treatment? Do we know anything about diagnosis and treatment in those same communities? Or is that just a blank on our evaluation map?

**0:41:28 HK:** Right. So, we do... The Malawi distribution, I mean, we do follow the diagnosis because it's relevant to malaria cases. And we want to see malaria cases before and after. So we know, for example, that a while ago, they were doing diagnosis in kind of a less preferred way and now they're doing it in a more preferred way, which is using these rapid diagnostic tests. But we're not tracking the whole effort because that's not how we see our role. What we're looking for is not giant problems that our donors can solve by themselves. We're looking for areas in which it makes sense for our donors to play a role, areas in which they can have kind of a localized positive impact and we don't expect them to take the whole thing on by themselves.

**0:42:15 EH:** There's always this challenge that you face in any program that you implement, which is that there's limited resources and that there's great need. So what we're trying to find is particular opportunities where the resources that we can bring to bear have a large, long-term impact. And that seems to be the case, we believe that that's the case with AMF. There's certainly other things that could be done and we'd want to evaluate them similarly.

**0:42:44 HK:** And there's a lot of organizations doing really important work that does not fit in the framework we're discussing tonight. And we're very aware of that and then I'm gonna talk in a bit about how we think about broadening GiveWell's scope, which will necessarily get involved in other health efforts. But it's certainly something to keep in mind that we're looking for this area where we can help and not making any claims that the other work going on is not valuable.

**0:43:11 S?:** On the case of GiveWell, will you say that there might be possible side effects like create jealousy in the neighborhood?

**0:43:24 EH:** So, I think the question's about GiveDirectly where they're giving out cash and they're going to neighborhoods and just giving people money and so, there's this question of what are the

unintended consequences or possible harms being caused because they do this? And, similarly, this is a question that applies in everything that charities do whether that's potentially destroying local businesses and there's all sorts of examples, diverting skilled labor from other activities to the charitable activities that might be less beneficial than what people otherwise would have done.

**0:43:56 EH:** So, in every case where we look at, a program or an organization, this is a question that we ask. And if you go to our website and you look at the reviews of organizations, there's always this section about negative or offsetting impact where we essentially try to raise every question we can about what could be going wrong and say what we know about it. In the case of GiveDirectly specifically, there's a couple of ways that harm could be caused. So, one thing is you could distort local markets, you could create inflation, you could have jealousy among the people who receive money, relative to the people who don't receive money. You could have people spending money on things that are really harmful like alcohol.

**0:44:35 EH:** I mean, not really harmful markets but certainly not... Don't seem as beneficial as some of the types of investments that Holden talked about like the high rates of return. The independent academic evidence of this cash transfer programs don't find these effects when they look for them, so they don't find the facts of increased inflation. They don't find evidence of significantly increased spending on alcohol or tobacco. GiveDirectly itself is also concerned about these things, and so, when they give money to people, they'll ask them whether they saw any evidence in the community of any discontent about who received money. And one of the things that... I mean, this might sound strange, but that we really like about GiveDirectly is they did find some problems.

**0:45:23 EH:** About a third of the people who received money reported that someone in the community was upset about the money being given in a certain way. I was in Kenya about a month ago, visiting them, before we recommended them, and I talked to people who didn't receive money. They didn't receive money because they were wealthier, and they said, "We're really poor, too." And, frankly, they were extremely poor and they said, "Why are we not given money? We had to take out a loan to do X, Y and Z, and we pay these high rates of interests. And here's these other people who are just receiving cash."

**0:45:54 EH:** But I think one of the things that... So they definitely measure these harms or possible harms. I think that on balance, they're not very large. No one reported violence, they reported some discontent. One of things that I liked about GiveDirectly is we've talked about this with them, and we've said, "Well, do you think this approach of giving out cash to some people in the village because they seem like they're poor, and other people who are, they don't just say, seem poor meet this criteria that they've developed for particularly poor. Giving them money relative to other people who are extremely poor but just less so, I mean, is this really a good idea or just this possibly cause more harm than good?"

**0:46:34 EH:** And they say, "You know, this is a good question. It's something we thought about and it's something that we're ready to gather the information so that we can do things differently in the future, if that is what makes sense." And with a lot of the other organizations that we've interacted with over the years, it's like these questions that are so important to ask about whether the program is accomplishing good or accomplishing more harm than good or not even being asked, and the organization's not sensitive enough to what could be happening so that they can adjust course and do things better.

**0:47:07 S?:** So you said that it might be surprising that you were actually glad to see some of the discontent. What was the reason that you were actually glad to see some of that?

**0:47:14 EH:** So, the question is, why would I be glad to see discontent. Because too many times when we've asked charities for information, all you get is positive information back. And it's just... Doesn't anyone who's tried to run an enterprise that is complicated and try to solve a difficult problem, you'll always gonna run into problems. Things are always gonna go wrong and it just doesn't pass the commonsense test that that's how things would be. And so the fact that they found some discontent, indicated to us that their evaluation method was working at least at a very basic level of...

[background conversation]

**0:47:48 HK:** Or that we are reading data and not PR in a sense.

**0:47:52 EH:** And one of the ways we often talk about evaluation as a good evaluation is one where the organization could ultimately demonstrate that it was not worth it. And if you read an evaluation and you realize that there's no possible way that this report could have come out negatively, well that's not an evaluation, that's marketing material.

**0:48:08 HK:** And this is a recurring theme with GiveWell. I think a lot of the things that make us unique are our willingness to talk about what's hard and what hasn't gone well and, you know, we've done it ourselves. I mean we have a page... A link on our front page that says the mistakes that we've made as an organization. Also when we make a recommendation, we publish regular updates and some of those updates have bad news and we don't hide it, we don't sugarcoat it. This is something that we think the world of giving badly needs because if all you can talk about is the positive and reward people for giving, you know, certainly it encourages people to give but it also makes it really hard to have kind of an honest, substantive, intellectual discussion. And so we're trying to move past that for the set of individual donors who are interested in doing so and we think it will make them better givers.

**0:48:52 EH:** So if what you're looking for is the great story where everything goes perfectly and there's no problems, then GiveWell is not the right place for you. And if you want to hear all of the problems that we've found with the best charities we've ever come across then you've come to the right place.

[laughter]

**0:49:07 HK:** If you're out of context and you read the page reviewing our top charities, it will look like we're just ripping on them. I mean, you know, by the standards of what you would normally read about a charity which has always been kind of a communications issue for us, but we want to really be out there with the strengths and the weaknesses.

**0:49:26 S?:** Having worked with a number of charities and things like this, what would have to happen for interested charities to be able to publish the data that you guys are finding. So is it an issue of charities don't have the funding to do the studies?

**0:49:52 EH:** So the question is what would have to happen for charities to be gathering and sharing the type of information we look for more consistently? I think the... I'll give a really simple answer, and Holden you can add if you want. I think the simplest answer is that right now, the incentives just don't exist for charities to do this, and that's the biggest thing that I think has to change and that starts with donors. Donor dollars are out there looking for evidence of impact or whatever compelling case can be made for impact along with careful, critical evaluation of a program, then I think charities will follow because they'll be forced to. But on the flipside, if donor dollars continue to follow the paths they have in the past, which is rewarding great stories, and great marketing, and great fundraising events that is just necessarily not going to incentivize charities to do the type of work that they need to do.

**0:50:54 EH:** Another thing that needs to change for donors if they demand it is they have to pay for it. Sometimes donors get so focused on low overhead, which is, I think, a pretty irrelevant metric and in a lot of ways a very harmful one. You know, donors say, "We don't want to see you spend on an evaluation, we want to see you just give more to the children," well obviously charities are not going to be in a position to do that. And so the bottom line is the funding for charities comes from donors and donors need to care enough about this information to demand it but also pay for it.

**0:51:22 HK:** So, I agree with that, kind of in a broad level, but I do think that it's also worth mentioning it's possible that charities are often not going to be able to make things happen by themselves. So if one charity says "This is what we want," we've talked about having multiple studies and having the studies be of a certain quality. Usually the studies that are high enough quality are coming from academic research and so there's a whole lot of moving parts and a whole lot of things that have to change for an intervention to have a strong evidence base behind it. So I think the road is kind of long and twisting, but certainly I think it definitely begins, and in a big picture way, it's driven by demand and the more money is kind of chasing the things that are documented to work, well the more work will go into documenting them.

**0:52:11 S?:** So have you seen positive effect from your technique of going to charities to say, "You could do X, Y, Z better," or "We can't recommend you because you haven't held this study," or you seem... Or getting positive message. "Oh that's a good idea, we should do this study", and this and that.

**0:52:29 HK:** Yeah. So the question is do we get positive feedback when we kind of critique a charity and say what it should do differently and then do they go and do it differently? So a couple of answers, I mean, one... Like I said, it's often there's a lot of moving parts and a lot of things that have to happen and so because of that and because of our intense focus on top charities and the amount of work we have to do to investigate even one charity, we're normally not focusing our efforts on critiquing non-top charities. It's just it's not something we really specialize in and for that reason it's kind of hard to answer your question. There is the occasional charity that says, "We love you guys and we want to do things your ways. Could you give us some advice?" Usually we will give them some thought but it's hard to tell if they're even sincere in that or if they're trying to make a good impression.

**0:53:18 HK:** This is kind of a common problem to have is kind of a funder or someone who's recommending actions from funders, and it's not something that we put a lot of effort into. But what we do do is, certainly with our top charities, we're in much more intense conversations with them about what is an issue for us and what we'd like to see done differently and none of it is



prescriptive, none of it is, "We will recommend you conditional on such and such a change." And the gifts we recommend are always unrestricted. They're always no strings attached in this category. But certainly, I mean, with AMF, they used to be our number 3 charity and we said, "We're concerned about kind of long-term usage data, because the nets are distributed but we aren't seeing anything about whether they're being used. It would even be nicer to have data on malaria case rates before and after." These are changes that AMF implemented. Now, is that because of us or is that because AMF really just wanted this data? I think it's a combination.

[background conversation]

**0:54:25 HK:** Do you want to talk about that? Start of education? Do you want to?

[background conversation]

**0:54:33 HK:** Okay.

**0:54:33 S?:** Perfect, we need to know.

**0:54:36 HK:** Okay.

**0:54:36 EH:** Yeah?

**0:54:37 S?:** A couple questions. Number one, how do you get folks interested in the areas that you guys have dived into, because some donors may be interested in education, health and things like that? And secondly, how do you build trust in donors that the research is sound, that your recommendations are good compared to something like Charity Navigator?

**0:54:57 HK:** Right.

**0:54:57 EH:** So the question is how... We cover certain areas. How do we develop the interest among donors in those areas, and then also I guess, why do donors rely on us or how do we develop the donors' trust in us relative to other options that are out there? The approach we've always taken to our research is trying to find the places that we would give our own money. Then allowing donors that have similar values and similar interests to rely on that research to give. So there are donors out there that we meet. I mean, many of them who say, "I am really interested in homelessness in St. Louis," or "Women's Immigration in Boston." And in those cases, we are not going to be able to serve that donor well. That's not the person we're trying to serve. That's not the market we serve.

**0:55:47 EH:** Our value added is trying to do the research that's going to be most leverage because it's going to apply to the broadest group of donors. And we think that that is going to happen when we take a, what we call, issue-agnostic approach to giving, which is not focused on any particular issue, but rather trying to look across issues to find the ones where we think dollars could do the most good, and then focusing in on them. On the trust side, how donors decide to rely on us. The other evaluators that are out there, the Charity Navigators of the world, are really doing something very different than what we're doing. They're trying to apply a rating to many organizations on objective criteria which largely is very different than the types of questions you'd need to ask to know whether a program is accomplishing good or not. So they're looking at things like, how much

does this organization spend on overhead. That's not going to answer questions like, well, are these wells still working two years later, which is the type of question that we want the answer to.

**0:56:53 EH:** In terms of the product we offer is very different. Then in the approach what we do is try to just put every single thing out there that we do so people can follow along with us, either by reading our research in depth on our website or following along with the blog and engaging in the conversation so people can gain an understanding where we're coming from and why we're reaching the conclusions that we are. And I think that most of the donors, who have supported us throughout our history, are ones who are following along with us closely enough that they've formed the view that we're doing work that's worth relying on.

**0:57:28 HK:** But just to be clear. That doesn't mean that they read every word that we've written. Probably no one's ever done that. I hope nobody's ever done that. I think there's one guy who did, but we hired him.

[laughter]

**0:57:42 HK:** You can spot-check GiveWell. We're trying to put everything we do out there so you can come, you can dig on some claim that interests you. Or you can read the blog and you can see how we're thinking. That is how a lot of donors develop trust in GiveWell, is they engage in some of what we do. They come to trust our thought process. They might come to an event, they might talk to us. And so it's basically like we want to put it all out there, so that any part of it that you want to dig on and check out you can do. And you can also know that other people are doing the same things. I forget what the first thing you asked because I had something to add on that, but...

**0:58:15 S?:** How do you get people interested in your area?

**0:58:17 HK:** Oh, yeah. One thing I wanted to just say about the long-term vision is that...

**0:58:22 EH:** While doing that, you can also talk about the future research now.

**0:58:24 HK:** Sure, sure. This a good time to get into that. Okay. One thing about the long-term vision of GiveWell is that we know we're making judgment calls, and we know that we're giving where we would give, which is not the one objective truth where everyone should give. Our kind of dream, or our kind of vision for the future, it's not that everyone in the world uses GiveWell's research to give. We're publishing information about how much money we think our top charities can absorb. And if there were more than that, we'd take down there recommendations. But what we do want to see what we don't see today, and would be very different, is a world in which there's just a discussion about it. So we are trying to go out there and say this is where we're giving and this is why, and have other people think about that and have people who have less time to devote than us maybe follow it, but also think about it and debate it, and do that kind of work for themselves. And so that's kind of the big picture.

**0:59:21 HK:** So, yes, this is probably a good time for me to talk about the future direction of GiveWell's research. We are going to continue doing the work on the charities we've talked about. The strong evidence, the direct translation of dollars into people helped, and that's where we started just because that's where... I mean that's the logical place to start, in my opinion, when you're trying to give and do a lot of good, because in a way it's the easiest. And it's not been easy, but I think it's

easier than the other kinds of giving. But the world of philanthropy is much bigger than that, and eventually we need to get into areas that don't work that way and that are very high risk and very high reward when they work out.

**1:00:02 HK:** So, one of our goals for the next year is gonna be to really make traction on that, and we've done a fair amount of work on it this year. And we expect it to be a long process. So, we're kind of trying to broaden our scope and look at the whole world of giving opportunities rather than just the things we've described, but it could be a long time before this results in recommendations. And the reason for that is that we've just found repeatedly that it's very hard to make a good recommendation without having a fair amount of context. Ideas that seem really exciting the first time you hear about them, once you start investigating them they often seem very, very different and are often less exciting. And a lot of things that don't sound so great at first sound more exciting when you have more context.

**1:00:45 HK:** And so, we're aware that it took us a long time in international aid and in this direct aid model to get to the point where we felt pretty good about our recommendations. And we expect it'll take a long time, too with these other causes. So with that said, here are some of the things we're doing to try to make progress on attacking the larger world of giving opportunities and being more open to anything. One thing we're doing is we're working with Carrie from Good Ventures on talking to some of the biggest foundations -- the Gates foundation, the Hewlett foundation -- and trying to understand as much as we can about what they think the best giving opportunities are and why they think that. And just understanding as much as we can their thought process and their operations because they've been doing this a long time and we think we can learn a lot from them.

**1:01:36 HK:** Another thing that we're doing though, is what we call strategic call selection which is trying to identify broad areas of philanthropy that don't get enough attention from existing funders, for how promising they are. And so, if you go to a foundation's website, they usually have focus areas and it makes sense because you want to have context in the field and you'll see that a foundation works on education or environment, and a lot of times when you ask how they chose them its kind of, "Well, this is what I was personally interested in, so we started there." And this may be the best way to give, but we also think another possibility is that you might, by broadly surveying things, by analyzing the data about how foundations spend their money, by reading as much as you can find about the history of philanthropy, what's worked? What hasn't? And actually, we've had trouble finding much on that, but we've read what we could find.

**1:02:25 HK:** You might be able to start identifying, "Okay, this area is promising, but there is barely anyone working in it," or there's not as many people working in it as in these less promising areas. So, trying to find the causes that look especially promising to dig in on is another thing we've been working on. And then a third broad category of what we are doing, is we do want to go deeper on some particular causes that we think are particularly important to our future learning. A big one is scientific research and especially biomedical research, so a lot of giving is direct aid which we've already done a lot of work on. A lot of giving also is things like cancer research and other forms of biomedical research, and this is something we know very littl

e

about. We expect it will take us a while to feel comfortable with it and have enough context to make a recommendation. We think it's very important that we get started now.

**1:03:16 HK:** Another area that we've looked into, kind of with a little bit more depth, is kind of an unusual one. And it's an area that we call meta-research. So, I mentioned before how we've run into this problem through our current work of a lot of studies just not really seeming optimized to do the most good for the world. For example, a study that doesn't share its data and code, there could be good reasons for this from an academics point of view. There's real incentives for the academics to behave the way they do, and we're not accusing anyone of doing anything wrong. But, in a world where the academics were kind of... Their incentives were aligned with good for humanity, we do believe you'd see a lot more data and code sharing than you do. And we think you'd see more per-analysis plans, which is something we mentioned earlier. And, there is a lot of other things we think you'd see.

**1:04:05 HK:** And so, we think this is kind of an interesting idea. That there is this world of academic research. It is huge. It has a lot of money in it. It has a lot really brilliant people in it. It does a lot of really essential work, but a lot of it is kind of people following their own incentives and playing by rules that are different from do the most good for humanity. And so, are there nudges we could make? Are there projects we can find that push things into a different equilibrium and align incentives more? That could be a very leveraged impact, and we don't know of any other foundations that have been working on this for a substantial period of time. We know two other major funders who are just getting started like we are, and so, we're working with them on this kind of area and asking around about it.

**1:04:50 HK:** So very broadly speaking, the work we've been doing on this area, in this kind of newer space, is similar to what a lot of people do when they just start out as a major philanthropist. They're trying to get some basic context. They're trying to learn basic things, talk to people, and get ready to form a strategy. But, the big difference is coming back to transparency, is that we have a mission of sharing what we learn because we're not an organization that exists to give one person's money away or an organization that exists to make public recommendations and start a public discussion.

**1:05:25 HK:** And so, this is what it comes down to in a lot of important domains, politics and business, there's a lot of public analysis and a lot of public discussion and argumentation about the best way to do things. And yet, public discussion of how to give is pretty rare except when it's very superficial and when it's all about encouraging giving and saying nice things. And that just doesn't seem right to us, and there is substantive intellectual discussion about giving, but a lot of it is happening behind closed doors. So, we don't think that needs to be restricted to the world's wealthiest people and biggest foundations and philanthropists and we think that could be a global public conversation about how to do the most good and that's kind of where we're trying to bring things and that's what we think will lead to the best outcomes overall.

**1:06:15 HK:** So we'll take some more questions now, questions about the future direction stuff that I just talked about, those are welcome, also if you guys have more questions about top charities, we're happy to field those, too, and we're open to anything.

**1:06:30 S?:** I'm interested in iodine supplementation... What are your current feelings about using some of your top recommended charities, that one could imagine being very efficacious, there's some evidence that very little iodine increases the IQ of kids. What do you think about this?

**1:06:47 EH:** Yeah. So the question is about... Basically a project... Program called salt iodization which is fortifying salt with iodine and the evidence is pretty strong that when you iodize salt, you accomplish a lot of good in terms of people's long-term cognitive development. So, this is a program that we did look into at great depth this year, in particular, trying to find an organization that we could give to where we believe that by giving to that organization, you would accomplish the outcome of more people utilizing iodized salt and we were, frankly, not able to find that. There are a few major players who work on this program.

**1:07:32 EH:** What we found is that in many cases, there are so many interconnecting parts between the salt producers, the government that needs to get onboard, the funders that need to provide some of the funding, that the obstacles to doing more salt iodization was not a funding obstacle but a coordination, negotiation and agreement obstacle among all of the players. And so we have not found an opportunity to which we could direct funding that would lead to more salt iodization, though it remains one that we're really interested in, and if we were able to find that, it's one that we would investigate very thoroughly.

**1:08:12 S?:** And sort of following on that, are there other things like salt iodization with strong evidence but there aren't necessarily organizations?

**1:08:18 EH:** Yeah. So the question is whether there's other things like salt iodization, and have that strong evidence but questionable ability to use funds to do more of them, at least in a relatively direct way. So, I mean, we talked about immunizations before, and that was a big area that we also spent a lot of time on, in this year in particular, trying to find a way to give, and there's a long write-up on that on our website. Another program that we're kind of still in the midst of looking at is a nutritional program called Vitamin A supplementation where they just give kids Vitamin A twice a year and the evidence is pretty strong that you reduce mortality by giving out these supplements.

**1:08:57 EH:** That's one where we've gotten a lot of like differing points of view about whether the gap exists such that you can provide funding and accomplish more Vitamin A supplementation, and so we're still investigating that one. But basically with all of the programs we look at, there's these different questions of whether there's a gap for funding, whether the program has strong evidence and we're trying to triangulate the ones that seem to have really strong evidence and we just need to find that gap. And it can often be difficult to answer because when you go to an organization, you say, "Well, do you need more money to do more?" They'll often say, "Yes," and so the challenge for us is finding a way to evaluate their answers to those questions in a way that we're able to hold them accountable after the fact, to make sure that they've told us something that is true. Yeah?

**1:09:47 S?:** So I have been a fan for a long time but part of it came from when you kind of mentioned earlier that there's sort of this basic disconnect in some ways in the non-profit world where non-profits are trying to raise money from people who are not their customers, so that creates a natural disconnect that by sort of part of research and actively looking at the organizations that you kind of reveal that and hopefully refocus them in right direction. And so I'm a little bit, I guess, worried about the potential shift toward funding research which seems like a very different kind of market and I don't know enough about it to say that you couldn't improve it, but I'm wondering what makes

you think you can?

**1:10:31 HK:** Sure. So that the kind of comment and questions, Gordon's been a fan for awhile because of this model where people who are serving... Charity serve people who are not their customer. So, you give to a charity, they go help someone else and kind of there's this broken feedback loop where things don't work the way they're supposed to and I guess the question is how does this apply to research, is that right?

**1:10:51 S?:** Yeah. What makes you think you're gonna help people decide how to cure cancer better than whatever way that they're doing it today?

**1:11:00 HK:** Right. So, certainly I mean, we're not... I think that the broad thing we're looking for is giving opportunities and specifically non-profit giving opportunities, so that would be anything where kind of there's no... Very broadly speaking, there's no constituency for a change other than potentially like philanthropists and non-profits and there's no one who's going to fight for this thing to happen because it benefits them directly. There's kind of no interest group. And that's kind of a broader way to put what you said, that the people who are helped by bed net distribution, they're not gonna make bed net distribution happen because they kind of don't have the power and influence and money.

**1:11:38 HK:** And I think research actually fits the same description in that broader sense, that the people who were helped by research are extremely dispersed and so, you kind of have the same broken system where you have... I might give to a cancer organization but it's not like I'm... If they cure cancer, it helps the whole world and it's just that and it's this very dispersed constituency where you really have some people trying to decide on behalf of other people what's gonna happen. And another issue with research is the fact and it's so hard to measure and hold accountable. I mean it's not like going to a restaurant where I pay and then I see how I felt about the food and if I want to go back, I mean you put some the money, you give some money to a prestigious university and then we continue not to have a cure for cancer and now what?

**1:12:26 HK:** It's very hard. Research is so risky that it's hard to attribute what happened to that one thing. So to me, it's broadly the same kind of brokenness where the only people who are... It's these kind of altruistic actors are the only ones who really have the incentives to get it right. And I don't think that all of the incentives in academia are altruistic. I think some of them are. But I think there is some places where academics are behaving differently from if they were just trying to create maximum good, and someone who's coming in with the kind of altruistic goals can say, "Alright, it would be better if things were done differently, what can we do about that?"

**1:13:03 HK:** And I think it's also worth observing that when you look at the things that philanthropy has done well and if you're trying to list the big success stories of philanthropy in giving, I really think research has the most impressive ones. So there's a lot of like cases where there is a philanthropist who got involved in funding scientific research and something happened that wouldn't have happened otherwise, it's really impressive and I think that it matches this idea that it's not clear who else is going to do this. Certainly the government does it, but the government to me it kind of acts like a philanthropist, it's there's no particular constituency that's kind of holding the beacon of fire.

[background conversation]

**1:13:50 HK:** Yeah. So this is a good example, I just got reminded of an example, the Cochrane Collaboration is a group whose work we've relied on. So what they do is they take lots of studies and they kind of meta-analyze them, which means they're digging up all the data and trying to look at what the accumulative effect is. And with the bed nets, there's 22 randomized control trials, we couldn't review each of those studies. To review each of those studies and pick them apart and look for the overall observations takes an enormous amount of effort that Cochrane Collaboration put in. And there's a Cochrane review that we relied on, because we've checked out Cochrane and we think it can be relied on and we've written about why that is.

**1:14:28 HK:** So, Cochrane is an interesting example because we have... It's one of these things that made sense from an altruism point of view, like, yeah, we should have someone collecting all these studies and putting them together to get the answer. And in some cases and some countries, they have a constituency, which is the single care and health program. And so the government is paying for medical care and they need this information or they want this information because they're spending money and they want to know what works. But in the US, it's like we did and we talked to the US Cochrane Center and they had literally no revenue coming in and we talked to various funders in the US who said, "Yeah, this is a cool sounding group, they do cool work, but it's just not one of our issue areas, it just kind of picked our issue areas." So that's just kind of an example of the dilemma I'm describing.

**1:15:14 EH:** Yeah?

**1:15:16 S?:** Do you think GiveWell would be more effective organization if you told great stories and had case studies about your research and show people that's helping with bed nets. And given that these kind of great stories in charity probably won't go anywhere. This is how people think and you talk to people and influence people. What's your communication strategy for starting this discourse? So a lot of your future direction is about starting the whole discourse about raising awareness about public giving. What's your communications strategy and how do you keep that evidence base while still recognizing how people actually think about it?

**1:16:04 EH:** Right. So the question is what is GiveWell's strategy for spreading the word about what GiveWell is doing as effectively as possible? This is a big goal of ours. We want to increase our influence significantly in the future. We're happy with the progress that we have made thus far in it's... I think as our influence has grown over the last few years, we're... It's been like strong, positive, fast growth, but it's still not something that we're satisfied with and we want to grow a lot faster. And so, there's a few things I'd say to you. I mean first of all, on this question of should we be using the tactics, so to speak, that other charities use to convince owners to give? On the one hand, I would guess that that's not going to be our particular value added. I mean we're researchers, we're interested in the data, and we're unlikely to be better at, let's say than World Vision or Save the Children, at putting together great videos that really pull on those heart strings to get you to give.

**1:17:12 EH:** Also there's this difficult tension, because while we want to communicate effectively to people and get them excited about what we're doing, if we just fall back on these same old story telling that other organizations use, well, in a lot of ways, we're falling into the same trap and we're not quite accomplishing the goal that we want of encouraging a more thoughtful conversation about giving. The final thing to say on this point is just that, where we sit right now with GiveWell. Let's just say there's these two competing potential ways we can spend our time, outreach and research,

and we still see... Even though GiveWell is still relatively small, we still see research as the more pressing priority to which we'll allocate our time. And we still see our time as really the most precious resource we have to allocate.

**1:18:08 EH:** And so when I think about GiveWell two years from now, well, there's one world in which we've put a lot more work into marketing and we've grown faster than we otherwise would, but we still are sitting here with the same level of understanding of different causes that we talked about tonight. And then there's another world where we continue to grow roughly the pace that we have but two years from now, our understanding of the research we do is far better than it is today. And that, you know, I don't like either trade-off but I'm willing to accept the slightly slower growth on marketing for the fast growth on research and not vice-versa.

**1:18:42 HK:** And the quick way... What?

**1:18:45 S?:** I was gonna say as a trade-off short term for better research but there would be less people that are reached.

**1:18:52 HK:** Yeah, less people that are reached.

**1:18:55 HK:** Right. I mean, a quick way that I'd put it is that we're still... At this point, I think we're still refining the product very much so. And the product is our research and our recommendations. And to the extent that we're doing marketing we're kind of we're reaching the early adopters. So people who don't need to be kind of reached out to in this way that centered around videos and emotional stories, that people who are actually really excited to do something different, who instead of coming to a black tie gala are like here eating pizza and having honest conversations about giving.

**1:19:26 HK:** There are enough of those people that we feel that made sense to target them for now and we're learning a lot by doing that and we're finding our research a lot by doing it. And there may come a day where we feel that our research is not as improving as fast, it doesn't need to improve as fast, and our growth is slow too and so it's time to focus on the growth. But I don't think we're at that point right now and I think ultimately we're much better served for reaching our current audience by really maximizing our command of the research which I think still has a long way to go.

**1:19:58 S?:** One of the things you've mentioned before was the Cochrane Collaboration. What are your current thoughts on whether it would be recommended in the future?

**1:20:17 HK:** So the question is, a couple of topics we've talked about in the blog this year, one was the Cochrane Collaboration which I talked about. They're the ones that summarized medical literature. The other one is Global Catastrophic Risks, so the kind of themes that kind of... This is another example of funding without a constituency because it seems like it happened like climate change and other things as well but would be such a large disaster and that there's kind of no particular constituency feeding it, to lobby for what we need to do. So what are the chances that we recommend one of those in the future?

**1:20:50 HK:** I think it's really hard to say. I think that the plan I outlined is the plan that we're most likely to follow which is kind of the causes we're almost likely to be go deep dives on are



biomedical research and this meta research thing, both of which are related to us needing to understand more about the general role of funding research which to me is like one of the three big areas of giving. I mean, most giving is either direct aid, which we've already done a lot of work on or its funding research or it's political advocacy which is kind of maybe last on our list and the one that we would...

**1:21:23 HK:** I think we're gonna be putting some effort into learning more about it in a very preliminary way, but we'd go after scientific research for us. So that's kind of the plan and it can be a long time before we make recommendations, so I don't really know what to say about that. Cochrane does fit pretty well into this kind of meta research within the biomedical world and I think there's a good chance we'll be looking more at them. I think Global Catastrophic Risks, we have done some work on them, we've written some on them and we will be writing more but I think it's fairly unlikely that we get all the way to a recommendation with one of those last year.

**1:21:57 S?:** Have you thought of giving recommendations that are tiered based on the size of giving. I have a \$100 people might be attracted to that.

**1:22:05 HK:** Right.

**1:22:06 S?:** A 100 thousand your top three might be attractive or I have a 100 million, that's probably not one of your top groups.

**1:22:12 HK:** Right. Yeah.

**1:22:14 EH:** So the question is whether we think about giving... Whether we think about providing recommendations based on the amount that a donor will give. So there's one consideration that you raised is whether the amount that someone would give would overwhelm the ability of the organizations that we recommended to utilize that funding well. And so this is something that we're aware of for all the organizations that we have. We have a view on how much funding they can take in if they take in that amount of money, we take them down and we say, "Look, we don't think that this organization is any longer should remain at the top of the list because they've received the amount of funding that they need to to accomplish what they can." And this has happened to us in the past where we had an organization that was our top recommendation and we removed them because they received sufficient funding. Up to that point where they would be overwhelmed, I think that the organizations we recommend are equally good for the \$100 versus the higher donor.

**1:23:14 S?:** No it wasn't actually being critical about it. It just, when we think about it, research on new rice varieties to improve yields, if you have Rockefeller money that makes a lot of sense. You should fund these, and you've mentioned this earlier and I've been thinking about it. If I have a \$100, maybe that's not really what I should be thinking about doing with my \$100. And so that's why I was thinking that depending on the amount one has, you might have just different criteria for selecting organizations.

**1:23:46 HK:** Yeah. I think that's very true. I think there are very different options that are available to a large philanthropist versus a small. And this is something that we learn more and more about as we talk to foundations, as we think about what we can do in certain causes. You know a large philanthropist can say, "Well, there's no organization doing expert, there should be, let's go pull together the relevant people and found this thing," and that takes a certain level of commitment and

that takes a certain amount of money and you really... There is nothing you can do with a \$100 that's any percentage of that.

**1:24:14 HK:** So I think in that sense, we're aware of this and I think it's a point right now our recommendations are targeted at people who as long as you're not gonna give so much you overwhelm the charity, we think it's a perfectly valid fit. These are kind of let we've talked about the easy ones and in the future we may say, "Look, these recommendations make sense for major philanthropists, it doesn't make sense for \$100." However, we also think there is an interesting fact that some of the smaller donors would like to act more like a larger donors. Some of the smaller donors might be interested, I know I would be, in the idea of pooling your funds and doing things that you couldn't do as one person giving \$1,000, and that's something that we've talked about, how do we experiment with, how to set that up so it's something that we would consider.

**1:24:57 S?:** And part of the new direction came from GiveWell Labs, which was a response to having the potential to influence larger donations.

**1:25:07 HK:** Right.

**1:25:07 S?:** So we may not be able to do that research with individual funding, but with larger donors...

**1:25:19 HK:** Yeah, that's right. Yeah, that's right. Going into the broader area of philanthropy is definitely partly response to see interest from larger donors where we feel that we should find options other than what we have.

**1:25:33 EH:** Yeah?

[background conversation]

**1:26:13 HK:** Right, so this is... The question is, are there... Is there an argument that larger donors because they have more options that they are sort of, and they have more time and more resources and more research, that they do a better job giving than smaller donors, and so smaller donors should give to larger to donors? And if that's the case, why don't you see the larger donors... Why isn't the Gates Foundation asking for money? Actually, they specifically asked not for... You know, they say, "Please, don't send money." Some people do anyway which is some kind of interesting. But... What is this dynamic? And I think it's an interesting question.

**1:26:46 HK:** So, one answer is that some large charities act like large donors in a sense. Right? They have their own staff, they have their own strategy, the problem from our perspective is that we can't see inside and like understand what they're doing to the point of having confidence in them. And if we could, we would support that. And so one way of thinking of Against Malaria Foundation is that it's a large donor. I mean, it's basically one person with an advisory board but it's a large donor than we understand well enough to say, "Okay, we wanna outsource this to that person and outsource to that donor."

**1:27:20 HK:** Now, that the private foundations cannot solicit donations. And I'm actually not exactly... I go to the universal answer to that, I think a lot of it is just like that would be distraction from their main mission. And processing the donations, doing the thank you's, the donor relations

could be a major distraction. So, it becomes a same issue with them but I think, conceptually, it makes a lot of sense to give your money to a larger donor that is able to do other things but you need to be able to see enough about the larger donor to have confidence in it. And in some ways that's kind of the role that GiveWell may start to play. I mean, with the broader philanthropy stuff, we may say, "Okay, this is what we would do as the larger donor and as a smaller donor you can contribute into a pool that will do it, if you want."

**1:28:07 EH:** Alright, we're going to take more question? Yeah.

[background conversation]

**1:28:46 HK:** Sure. So, the question is... Would we recommend giving to the Gates Foundation's charities and do we believe their analysis is correct and is good? So, we have spent some time talking with the Gates Foundation and not all of what we've discussed is stuff that we are willing to talk about publicly and that's because we always honor requests to keep things off the record and all that. What I will say is a couple of things. One, I think the Gates Foundation is a pretty impressive organization. They've been around for a while, they have a lot people on staff who seemed very knowledgeable and are very reasonable and who I found like pretty impressive. Another thing is that they are doing something very different from what we been talking about all night.

**1:29:32 HK:** They're often much more in the category of high risk, high reward. So a lot of the work they are doing is either kind of focused on advocacy like trying to pool in money from bigger donors and even bigger than them like government donors, into certain interventions. Or it's focused on kind of research and development like developing malaria vaccines, for example, or its focus on something like eradication, polio eradication these are kind of these very ambitious, very high risk, very high reward things, and so those are very different from what we focused on and in a way they're harder to assess.

**1:30:06 HK:** And the third thing I'd say is that it's not easy, especially because of the nature of what they're doing, to assess exactly what their analysis is. And so, it's very hard to talk to them for the limited amounts of time that we have and really understand all the things that have gone into their analysis and all the things they believe and how well grounded they are and assess how much we would trust them. So it's kind of an 'I don't know'. The final thing I would say is that Gates does a lot of project funding. So if you look on their website and you look at who they're giving to, a lot of it is a grant to group X to do Y. And so, that's why I would not...

**1:30:46 HK:** I mean they give to a lot of different groups and often with kind of a purpose or a tag. And so the one thing that I would disagree with is that. I would not feel comfortable just writing checks to those charities. I feel like, if they have their strategy, it may be a great strategy but that's different from just going on their website and looking through these hundreds of charities and writing checks for them. That's not something that I would feel comfortable doing. Do you have anything to add?

**1:31:14 S?:** No, I don't.

**1:31:15 EH:** Well, I think we're gonna wrap up. Holden and I are happy to stick around and just talk informally if anyone wants to. We really appreciate all of your coming out and I know that this may... Maybe it's more fun but maybe it's less fun than your normal black tie charity gala and we

GiveWell Research Event, San Francisco, December 3, 2012

really appreciate it because it's really important to us that people want to engage in this discussion with us because we can't do this alone, and it's only by everyone engaging with us that we're able to accomplish what we set out to. So, thank you.

[applause]

## **Thank You for choosing Scribie.com**

Cross-check this transcript against the audio quickly and efficiently using our online Integrated Editor. Please visit the following link and click the Check & Download button to start.

<https://scribie.com/files/db9f99732120a0c364b412cb032a01ec37a64c88>