

GiveWell research event, June 5, 2013

## **GiveWell research event, June 5, 2013**

*This transcript was compiled by an outside contractor, and GiveWell did not review it in full before publishing, so it is possible that parts of the audio were inaccurately transcribed. If you have questions about any part of this transcript, please review the original audio recording that was posted along with these notes.*

**0:00:00 Elie Hassenfeld:** Thanks everyone for coming. I'm Elie, I'm one of GiveWell's co-founders. This is Holden, another co-founder. Before getting started, I just want to introduce all the GiveWell folks who are here. So, this is Alexander, Eliza, who's over there, Tim, who's right there, and ER is here working with us part-time, and Cari, who works at Good Ventures and is on our Board. So, the basic plan for tonight is that we're just going to try to have an open conversation about GiveWell, and the way we're going to do that is that Holden is going to talk a little bit about what's going on in GiveWell Labs, which is the new areas of research that we're looking into. Alexander is then going to give a case study of one of the areas in GiveWell Labs that we've looked closely at, and then I'm going to give a little bit of an update on our more traditional work on developing world direct aid. In each of those sections, we're going to try to keep them relatively short and pause for questions, and we really just want to have an open conversation about what we're doing. We're going to talk about some specific things when we talk, but we're also happy to take questions on any of the different things we're doing. So, this sheet that I'm passing around now just lists out some of the major topics of work that we're doing, so you should feel to ask questions about any of them as we go.

**0:01:29 EH:** The whole thing's going to run about an hour and then we're happy to stick around and answer questions informally with anyone who wants to. One quick housekeeping note is, one of our core values is transparency and so we try to share as much information about our research process as we can, and so, to that end, we're recording this meeting. If there's anything that you say that you'd prefer not go into this recording that we're going to publish on our website, just say so when you talk or let me know afterwards or tomorrow. We're happy to cut things out. We just want to try to share as much as we can. So, before turning things over to Holden, I just want to give a quick overview of a big change that's going on at GiveWell right now.

**0:02:14 EH:** So, the big shift that we're going through is we're transitioning from just focusing on proven cost-effective international organizations to other areas of research. It's important to know that we're still committed to maintaining that high quality list of top charities that we've had before, so we continue to work on updating those organizations, researching new programs that could be similarly effective, and having an open door policy to any organizations that think they could meet the standards of our current top charities. We're really focusing a lot of our energy on this new area of research and there's a few reasons that we're moving into these new areas.

**0:02:55 EH:** One of the reasons is we are optimistic that by looking into areas that aren't constrained by our traditional criteria of evidence of effectiveness and cost effectiveness, we're likely to find programs that offer higher impact per dollar, where you can get more bang for your buck. Another reason that we want to look into these newer areas is that we think that they're the types of areas that are more attractive to major philanthropists. If you're an individual giving a thousand dollars a year, you have to give to an existing organization, but if you have more money to give, you have a lot more opportunities than the individual. And so, we want to look into those areas as well. And then, finally, I think the research that we've done in our traditional areas is in some ways, reaching some diminishing returns. We have focused on finding more or less the same type of organization for the six years that we've been in existence and we're no longer finding new

organizations that we hadn't heard about before or programs that we haven't assessed before. And so, the value of pouring more time and energy into those areas is lower, we think, than the other areas.

**0:04:04 EH:** One final thing to note is that in a lot of the work we're doing in GiveWell Labs we're working very closely with Cari from Good Ventures. We share an office, we participate in the research process together and in particular, a lot of this research wouldn't be possible if not for the... In particular, the access that Cari and Good Ventures helped us get to people who are excited to talk to us because we're working with Good Ventures on these new ventures of research for us. So with that, I just want to turn it over to Holden to talk a little bit about GiveWell Labs.

**0:04:40 HK:** Sure. So, I'm going to give a broad overview of GiveWell Labs and then talk about the specific stuff that I've been doing. And then, Alexander will later discuss some other aspects. So, we used to look for the best charities we could find and I think when we started GiveWell Labs, we were thinking of looking for the best projects we could find. So, where is a project that needs money to go forward that we think is really outstanding? And as we wrote in a recent blog post, we've shifted since then to really analyzing causes and thinking about our fundamental question being, which causes do we want to get involved in? And we wrote a lot of reasons about this. I won't go into great detail about why this shift has occurred, but a really illustrative example is if someone brings to you a project that is for a particular study on malaria control, it's very hard to assess whether this is a good project without having a lot of knowledge of what the different aspects of malaria control are, who the different players are, who the funders are, what they fund, what they don't fund, what the research already says, what's out there, and by the time you know all that, now, you're able to assess a lot of different malaria-related projects and it goes much better.

**0:05:47 HK:** And so, that's what the importance is of a cause to us. It's something where you need to get a lot of background knowledge and then you can evaluate a lot of things well. And, from a practical perspective, if you want to be good at understanding, giving opportunities, it seems like you have to pick some causes, you can't pick every cause, and you have to become really good at them, and you have to invest in them. That said, I think we're very interested in trying to choose the causes we focus on wisely and reflectively and systematically, and this is something we call "strategic cause selection", which is different from how we see other philanthropists doing things. I think it's common in philanthropy to say, "I'm passionate about cause X, so that's what I'm doing", and we're trying to look at all the plausible causes, gather a little information on all of them, then gather a little more information on others, and go from there, and pick eventually.

**0:06:38 HK:** So, one of the things we're doing that's pretty straight forward, is that we've got a long list of causes. And for some of them, we're gathering a tiny bit of information and for some of them we're gathering a little more and doing that kind of winnowing process. Alexander's going to talk about that. What I'm going to talk about is my efforts to kind of understand two very broad categories of philanthropy that apply to a lot of causes and that I think we basically have to understand in order to look at a lot of things.

**0:07:06 HK:** So basically, as I see it, most of the things philanthropists do today can be divided into one of three categories; there's direct aid, which I think we know a lot about and we've looked into a lot, and I think we can kind of tell you what the most proven direct aid is, and we have a lot of contexts, and it's kind of delivering things to people who can't afford them, for example. But the other two basic categories were things we know very little about. So it's common for philanthropists

to invest in policy advocacy, which is basically trying to get laws changed. And it's common for philanthropists to invest in scientific research, which is trying to develop fundamentally new technologies and ideas. And these are two things that our traditional work hasn't done and so, what I'm basically trying to do is get us grounded in those, get some basic context and get us equipped to ask the right questions for any given cause within political advocacy or within scientific research.

**0:07:57 HK:** So I'm going to very quickly go through them one at a time, what I've done so far at each one and what we're planning to do, but happy to take more questions at any particular aspect. So, first off, politics. I just want to open by saying that I think a lot of people are very skeptical and very kind of, have a negative emotional reaction to the very idea of getting involved in political advocacy and I can definitely empathize with it. I think that a few years ago, my view was kind of, "Well, on any given issue, there's a lot of smart people on both sides of the issue, there's a lot of well meaning people on both sides of the issue. If you get involved you're kind of liable to be overconfident and to fail, and to maybe do damage and all that."

**0:08:36 HK:** And sure, I still think these things are risks but I think the more I've learned about this area, the more I think that that's really not the right way of thinking about things at all and I think this is an area that's very worth exploring. So, what I would say is, I believe there's substantial points of consensus between people who have global humanitarian values. In other words, people who feel that people all over the world, all have value and they're interested in maximizing kind of the aggregate welfare in some sense, and if you throw in that people not only have those values, but are relatively informed, you'd get a lot points of consensus and a lot of the places in which policy diverges from what that consensus would say don't really have anything to do with global humanitarian values or with kind of distanced analysis.

**0:09:22 HK:** What I perceive the political world is involving is, there's a lot going on at any given time, there's a lot of bills that are up for discussion, there's a lot of discussion about what the next bill should be, and there's a lot of information to sort through, and when there's a well-organized and well-resourced lobby, what they're able to do is they're kind of able to scan that giant flood of information for the things that are relevant to their issue, and create passable bills that relate to their issue, and find things that are going on in debates that relate to their issue, build coalitions around their issue, create messages around their issue. And so a lot of it is sorting through and providing the right information to promote what they're doing, and this is the kind of thing that you need to be organized, and you need to be resourced to do this. And so, a lot of what happens is, you'll have corporations or industries, that or interest groups that have these well-resourced lobbies that they're able to push for their issues, and a lot of the things that might benefit humanity broadly can get lost in the shuffle or it can get actively defeated by that.

**0:10:25 HK:** And so I think a role for philanthropy has historically been and will in the future continue to be, find issues that matter a lot to humanity as a whole, but don't have that kind of natural constituency that creates a well-organized interest group, and then just create that interest group, provide the resources in the organization that don't arise so naturally, and that, I think, is pretty compelling. I think that we've definitely heard some supposed success stories, we've got a mission of understanding whether they're really success stories, that's a different story, but I think, definitely very plausible to me, that there are cases in which philanthropy kind of built up these coalitions around something like healthcare or immigration, that for a long time, we're getting allies and putting out ideas, and then when the moment came, they were really ready to take advantage of it, when their political moment came.

**0:11:14 HK:** So, it's a long term investment in building organizations, building capacity and then having kind of an unpredictable difference in how it's presented, and while it's very risky and while it's very long term, even a tiny impact on legislation, obviously, it's just an incredibly magnified effect. And so, there's a lot of people in philanthropy who believe this is just far and away, easily, the way to get the most effect for your dollar, and so therefore, if you pick the right causes, the most good for your dollar, and I think it would be a big mistake not to investigate it. What we're doing right now is, the basic framework we have is that we've got a list of issues that we think are important from a humanitarian perspective. So, an easy one to explain is, what is the size and strategy of the US foreign aid program and how humanitarian is it? How big is it?

**0:12:00 HK:** But we've got a lot of issues like this that we think are important and for each one, we're thinking about kind of trying to understand some basic parameters, like not only how important is this issue, but who are the interest groups? Who are the groups that are here? And is there space for something more? Is there something a philanthropist could do that isn't already being done? And so, I think we're getting ready to kind of go beyond the general understanding of policy level and get down to the individual issues and kind of try to investigate these, and look for a few candidates to get involved in and learn from as we go. So that's basically where we are in politics. Now, I'll talk quickly about...

**0:12:37 EH:** Do you want quick questions as you go or later on...

**0:12:39 HK:** Well, I think I'm going to try and blow right through this side of the research part and then I think we'll take questions on everything I just said, which I know is a lot, but, yeah. So, for scientific research, I think we're not as far along. We're starting with biomedical research which is very broadly, biology and medicine, it's about half of all research funding in the US. So it's a very, very big field, it's what a lot of philanthropy does, and things like trying to cure cancer, but also other diseases and other approaches. And just a couple of things that I believe in this area. One is that you need scientific advisers, it's just like, people are speaking another language and I actually can't understand what someone is claiming without the help of someone who's kind of literate in science, and so, we've been working informally with Dario, who's right here, who's a good personal friend and also, who I think has good broad scientific literacy and also is able to explain things well to me. We're starting to get going and we've been working with him, while at the same time, knowing that our choice of scientific advisers is a really crucial one and we have to get systematic about it over time and then put together the best possible team of advisers.

**0:13:46 HK:** And just to give a flavor of the kind of thing we're thinking about early on, I think there's a lot of different ways to slice scientific research. So, in politics, it's become pretty clear to me that you want to talk about issues, immigration reform, foreign aid. For each issue, there's people on one side, people on the other side and you look for the ones where you can get involved. Science, you could talk about which disease you should be funding. You could talk about, should we be working on cancer, heart disease, something else, but you can also slice it a lot of different ways. There are people who have that hypotheses that it's much more the stage of the process where money is really mis-allocated, that some of the most basic fundamental biology that doesn't have immediate application is maybe under-invested in, because of all the focus on diseases.

**0:14:26 HK:** Or maybe some of the later stage stuff that's kind of not as academically prestigious but that we need to get new treatments. So, that kind of hypothesis. And then there's a whole bunch

of different ways to slice the world and what we're doing at this moment, is we're trying to talk to the biggest, most prominent biomedical funders and understand how they see the world. How do they see the relevant way to slice it? What do they think is under-funded? What do they think the important debates are? And with Dario's help, that's what I'm doing and kind of simultaneously, thinking about who are adviser should be. So, that's kind of where we are on those two things, and with that, I'm just going to take questions on science, politics, GiveWell Labs in general.

**0:15:04 S?:** So in the politics side, you mentioned of finding things that don't have that sort of political lobbying group established, that might fight for this great thing that you've identified as being a great thing for humanity. Are you also sort of trying to catalog which causes have their special lobbying interest groups who are opponents of what you identified as the right thing?

**0:15:27 HK:** Yeah, yes.

**0:15:28 S?:** I feel like the problem is usually that there's a lobbying group that's causing the problem, rather than that there's the absence of a special interest group to fight for the right thing.

**0:15:38 HK:** Sure. So the question was, when we look at politics, I kind of talked about space for new philanthropists but are we also looking at who the opponents are and who's standing in the way of change, as opposed to who's already on the right side whatever we believe that is? And I would say, okay, definitely, we're looking into it. I mean, I think that's one of our questions and we're trying to broadly determine, what does the whole coalition of forces look like? And you do need to know that. I will say that one thing I've taken away from the reading and the conversations I've done, is that that side of things is maybe not as crucial as you might think. I think some of philanthropy success stories is just like the opponent looked completely insurmountable and they beat them anyway, in the end. And it took like a really long time but you have a couple of things going on. One is, the time horizons are really long.

**0:16:26 HK:** We're often talking about 10 or 20 years of investment before you can really expect to see a result because of this kind of coalition-building thing I'm describing. So, it's hard to predict, it's hard to look at today's environment and say, "Oh well, we can't pass a bill with this make up of people and so, forget it." It's hard to predict. And also, philanthropy does tend to enter into areas where kind of you could make a moral case for your side and so it's like once you get a little bit of traction, right? Things can really splinter in, and once the kind of powerful interests aren't assumed to be insurmountable anymore, all of a sudden, a lot of people would switch sides. So, I would say, yes, we're looking into it, but one of the conclusions I've come to is, you shouldn't run away from an issue just because the opposition looks strong. There are cases in the past where that would have been the wrong call. Yeah, Jake?

**0:17:18 S?:** You talked a little bit about... You were talking about strategic cause selection, how there's kind of return to each specialization, and there hasn't been something, like, GiveWell's been a very kind of generalist organization, survey and everything...

**0:17:29 HK:** Sure.

**0:17:31 S?:** But if there's so much return for a specialization, do you see GiveWell, in the future, maybe taking out a few issues and becoming kind of an organization that just does meta research or just does one or two things?

**0:17:44 HK:** Yeah. So the question is, GiveWell's historically been very kind of generalist, do we see ourselves becoming more specialists in maybe some areas? I would say we've actually become more specialists over time. I mean, we started off saying, "Whatever is proven, whatever is cost-effective, we'll do it." In the end, after several years, we concluded that having those criteria, you end up becoming global health people. And that's certainly a big cause, it's kind of a super cause but it's a cause, and it's what we know the most about. And I think in the future, yeah, we'll be specialized. I think our goal is to be smart about what to specialize in. And so what I envision is that we look across the whole landscape, we make some educated guesses about what areas to specialize in, and then we'll have ways of specializing in them. I mean, I think it's premature to get into how that's structured and what the staff do, but while we're specializing in them with part of our capacity, we're also continuing to ask that question, "Are these the right areas to specialize in, and what else should we specialize in?" I think that we aren't fundamentally generalists, we're fundamentally people who are considering all of the options in a systematic way, and making the best decisions that we can with the available information.

**0:18:53 S?:** One follow up question to that is, so maybe you specialize more in a certain area and there's a point in time where you've built in enough competency in that area, that you're better just doing that speciality, even though that there's other issues that you evaluate and say, "If GiveWell came along now, they would pick these other issues", but because we, at GiveWell, specialized and built up those competencies, it's actually better for GiveWell to keep focusing on that. So, there's kind of like a tension between those two things or something.

**0:19:27 HK:** Yeah. I mean, could we end up becoming stuck in an area in some sense, because our previous level of commitment raises the returns over what it would be for a new version of GiveWell? Yeah, I think that can happen. I think there's multiple reasons, that there's good reasons to stay in a cause for a long time, but I wouldn't quite think of it that way. I would say that we want to pick some causes, know that we're committing to them and reserve the capacity to go into new causes, and revisit those causes that we've committed to, and not stay committed to them just because. But yeah, if there's good reasons to stay committed, then it's a reasonable thing to do. I think that we never want to get to the point where we can't continue exploring and looking across, and taking some new incredibly promising thing and getting into that. I would hope that never happens, but if that turned out to be the way to do the most good, that's what we're doing, I guess.

**0:20:18 EH:** When we think about the focus now with the GiveWell Labs work, a lot of the focus is getting to the point where we're confident enough that we can commit to some degree of specialization in a few areas, not forever, but for a long period of time, where you have to be confident enough that it makes sense to put those resources in. 'Cause one of the things that we've learned from talking to existing funders -- and it's something we've learned from discussions with them, and seen -- is you just gain a lot by having that level of familiarity with the issues and the context, that you can make good decisions and distinguish good from bad ideas, or... Basically...

**0:21:03 S?:** Quick question. You said something at the beginning that sort of surprised me, that other philanthropists don't spend a lot of time on strategic cost selection. Is that even the case for some of the organizations that you have a lot of respect for, like the Gates Foundation?

**0:21:15 HK:** So, the question is, I said that other foundations don't seem to do strategic cause selection, is that even the case for the ones that we have respect for, like the Gates Foundation? It's

really hard to know for sure. I know that no other foundation is really using the same language we are, putting it in a way that's really obvious they're doing the same thing. There could be groups that are kind of implicitly or partially being strategic. As far as the Gates Foundation, specifically, I would say that I feel like they've done a really good job picking causes, I feel like that... Well, actually, I think within certain subparts of the Gates Foundation, they've done a good job picking causes, so I think the fact that they've really been on global health, I think within global health, they've made some decisions that seem pretty impact-oriented to me, but I haven't seen any sort of process like the process I'm describing now, on what GiveWell's doing. So, that's pretty much what I know.

**0:22:11 EH:** I mean, the other... Oh, what were you going to say?

**0:22:14 S3:** I was just going to chime in on the Gates Foundation. They started out putting computers in libraries, wasn't it? So, I think they got more strategic over time in selecting their causes, and certainly, we feel, in retrospect, that global health was an extremely strategic cause to select at the time that they were expanding to the size they're at now. But they didn't start out, I think, picking causes strategically. They started out with what Bill knew and what his passions were.

**0:22:44 EH:** Another piece of support for this notion that there's limited strategic cause selection is that when we've gone to funders and said, "Hey, what should we focus on? We're open to anything." They said, "Well, the first thing you have to do is pick the area that you are most passionate about." We said, "Well, we're not passionate about an area, we're passionate about doing the most good with the money that we have to direct." That conversation is so consistent, and it's so difficult to communicate that we're open to anything that does good, but I'd be very... At that level... And this is not only big funders, I think there's like some post on our blog about donors who talk to philanthropic advisers, not the biggest funders, but people are saying, "I'm giving \$10,000 to \$50,000 a year, maybe", and they have the exact same experience, where this notion of being open to anything is very outside of the scope of what people normally discuss.

**0:23:34 S3:** This has been my experience for a year, going around asking other philanthropists, foundation executives, philanthropic advisers, "What are the best things to be working on?"

**0:23:45 S?:** So, have you guys looked up the Copenhagen Consensus?

**0:23:47 HK:** Yeah. We're very familiar with...

**0:23:49 EH:** That's like the most similar thing...

**0:23:51 HK:** So, I think in terms of strategic causes...

**0:23:53 EH:** Say what the Copenhagen Consensus is.

**0:23:55 HK:** Sure, sure. So the question is, have we looked at the Copenhagen Consensus, and is that the closest thing to what we're talking about? So the Copenhagen Consensus is, basically it's a website you can go to. I mean, I don't know exactly what structurally it is, and they basically try to quantify the kind of payoff per dollar of different things you might do, like fighting climate change or improving the nutrition of children in the developing world, and they created a ranking of the



best causes. So that is certainly like the closest thing I can think of, to really trying to go through lots of causes and rate them and decide which one is the best. I mean, I will say that A, the set of causes they looked at is really constrained compared to the set of causes we're looking at. I don't think that they... I think it was generally a lot of developing world interventions, a lot of climate change interventions and not much else. I'm not entirely sure about that but I definitely remember seeing their list and feeling that way and looking at everything they considered.

**0:24:54 HK:** The bigger point to me is that we have a really fundamental difference of methodology with the Copenhagen Consensus, where their... I think this reflects, this is reflected in the causes they chose, that they're really insistent on quantifying everything, and I think that we, we have probably... I actually believe we have done the most thorough attempts in existence to quantify the good accomplished per dollar by a charity and I think along the way, we found and written up a lot of reasons why that approach is limited. It's kind of another conversation but the bottom line is that I don't look at their work and think that it's really usable for this purpose because... It's basically, it's opaque because of the way that they've done their quantification. And I've read their papers and like, a lot... It's like a very large percentage of a given Copenhagen Consensus paper is talking about which discount rate to use, because you change your discount rate. This is the rate at which you weigh the present versus the future benefits, and like everything changes, and it's that kind of thing that you can't really look at these rankings and like understand their reasons, or understand what's really good to invest in. And also, I think there wasn't much attention to what the actual giving opportunities are. I think it was all very hypothetical, it was all kind of model interventions based on studies so...

**0:26:09 S?:** My understanding is, they were not directing it so much in individual philanthropists as here in the UN or the governments, in orders of magnitude of dollars.

**0:26:18 HK:** Yeah. I mean, but I think we're also now directing ourselves to kind of larger pools of money, and you still...

**0:26:23 S?:** Well, the last thing...

**0:26:25 HK:** You still have to ask, I mean, "Who are the other offenders in this space? What are they doing? What are they missing? What are some things you can do that aren't already being done?" I mean, these were, these were much more kind of... If you could hypothetically replicate the results of the study, this what you would get. So, I mean, those are thoughts on the Copenhagen Consensus. I do agree it's the closest thing to kind of trying to write up what the most important causes are and I also feel comfortable saying that like, we have not found it useful in our research.

**0:26:53 S?:** Your description of yourselves as being "wide open" doesn't fit how I view you. So, maybe you can help me understand...

**0:27:02 HK:** Sure.

**0:27:02 S?:** I view you as focusing on what's easily measured, the dailies and not so much the quality of life, whether people are literate, whether they're free, unless that has an impact on disability-adjusted life years. So, what am I misunderstanding?

**0:27:21 HK:** Sure. So, the question is, is GiveWell kind of fundamentally at our mission or in our

core, are we focused, are we totally wide open to anything that might do good? Or are we focused on things that are measurable and quantifiable, and putting terms of dallies? So, the answer is the first. We're totally wide open. I would say historically, we've focused on things that are measurable. That was a practical decision. That was us saying, "We don't know where to start and we need to bite off something we can chew, and what we can do is look for proven cost-effective interventions." And for people who don't have the staff, the capacity, the context, the informal knowledge to look at other things, it makes a lot of sense to focus on what's quantified. So, I actually believe that the best giving opportunities I know today are our top charities, and we took the shortest path to find the best giving opportunities because I think when you know very little, finding things that can really be demonstrated in a formal way is the most efficient way to get the things that are doing a lot of good.

**0:28:18 HK:** But I think that now we have the plan and the ability, and the capacity and the demand. We're looking more broadly and I think that's always been the mission of GiveWell since day one, it's what our mission statement says, it's what we have always to do. Fundamentally, at our core, we're wide open and historically, there's been an emphasis on measurement but that's changing.

**0:28:39 EH:** Yeah. I mean even in our traditional, the traditional work that we've done for a long time, I mean there's a big role that quantification plays, but something we take flack from another group of people who are very focused on quantification is that qualitative judgments play a role in what we believe about the different organizations we rank. I mean, some organizations we think are more likely to solve future problems better than others, and that leads us to have a better view of them than others, and our ultimate goal is not to be... Our approach to this problem is not to say, "It's not quantifiable, therefore, throw it out." If it's playing a role in our thinking, we want to put it down and be transparent, and try to be totally open about all the factors going into our judgments.

**0:29:23 EH:** So I think we should pause on questions here and let Alexander talk a little bit about the case he will and then we'll have another round of questions and more time for the rest of anything later.

**0:29:34 Alexander Berger:** Sure. So, why don't I just start by giving a little bit more sort of context on how we're thinking about GiveWell Labs. So, Holden talked a lot about sort of the politics and science work that he's doing, but in addition to that, we're sort of doing two sets of work on more specific areas, and yet the basic idea is that we sort of envision ourselves going through a funnel over time. We have like a list somewhere of like a hundred causes that we think might be sort of competitors, so it's like the best cause type thing, and what we're trying to do is winnow that down so that in the future, when we're ready to commit some causes, we have a list of the ones that we think are likely to be the best bets. And so, what I'm working on is sort of working through that process and winnowing down that sort of giant list to the sort of manageable level, where you're going to want to make some specialized investments in human capital and whatever else. And so the sort of two levels in which we're working on this right now or sort of a medium depth level we're working on one investigation as sort of a test, and then a bunch of shallow investigations towards like just a very, very quick investigation of an area where we don't know much, to learn a little more so we can prioritize it better, going forward.

**0:30:42 AB:** And so I'm going to sort of briefly talk about one particular sort of shallow investigation that we've been working on, climate change, and it's one that we've worked on a little

bit over the past year or so, and it's not something that is... It's one of the ones we've done sort of deeper for a shallow level, but where there's tons of unanswered questions and where we don't know all that much. So, when we do one of these shallow investigations, we typically ask three questions as we're going through that process. So the first one is, this question like, "What's the problem and how big is it?" So, in this case, it'd be like, "How big of a problem is climate change?" The second questions like, "What kinds of intervention could you bring to the table? What can you do about this problem?"

**0:31:19 AB:** And then the third question we ask is, "Who are the current players and what are they doing?" Often, this takes like the very mechanical form of asking the question, like, "How much money is in the field already for a philanthropist?" That's like a really easy way to put that. And so the basic idea here is that you'll learn enough to help you get a picture of the fields, so you can better prioritize for future research about what the given opportunities might look like. So, I'm going to go through sort of like, there are preliminary answers to these three questions for climate change. So first, to try to get to the problem of how big of a problem is climate change, 'cause it's sort of a big tough question, we looked at the 2007 Intergovernmental Panel on Climate Change report, the IPCC as it's called, won the Nobel Prize with Al Gore, also a Nobel Prize winner then, and they're sort of like considered the foremost summarizers of the climate change research.

**0:32:11 AB:** And so we looked at their report, we sort of started with the summary. It just turns out that the summary had a lot of like really broad generic statements of the impact. They were hard for us to sort of interpret, not necessarily in dally terms but even in sort of like general humanitarian terms. So, let's say things like, climate change is expected to change the distribution of some disease vectors. And that's just like really hard to know what to make of, like how bad is it to change the distribution of the the disease vectors? And so we ended up taking a few attached from the summary that we felt like we can manageably go to and summarize to get a better sense of how that's supposed to be, and then we also looked at sort of econometric estimates of how that kind of thing is supposed to be. So, combining together in sort of economic or dollar terms, how bad all the damages are expected to be, and then putting that in sort of one final figure.

**0:32:57 AB:** And so, based on reading this initial summary, we decided to prioritize further work on, how many people are expected to suffer from hunger as a result of climate change? How many people are expected to suffer from water stress as a result of the climate change, flooding, extreme weather, things like hurricanes or especially hot or cold days, health impacts, biodiversity loss, and these sort of economic estimates that I've mentioned? And across all these areas where we went through, the sort of relevant chapters of the sort of like several thousand-paged book that they basically have put out, and what we found is that climate change is expected to have large negative impacts on pretty much all of these humanitarian dimensions and that sort of fits the kind of wisdom about climate change. It's going to be bad. That's sort of universally better if we did something about it.

**0:33:41 AB:** On the other hand, we also learned... And I think, I was at least a little bit surprised, when I learned this, which is that in a lot of cases, even if nothing's done, it actually looks like people are still going to be better off in the future because of economic growth, basically. And so if you look at projections for... Hunger is a good example. The number of people suffering of hunger, even if we don't do much about climate change, is still expected to fall in the future. Now, again, like that's not an argument against doing something about climate change. It's still way better. It's worth the investment to do something about climate change, but the actual like, if you're worried

about the potential downside or disaster from climate change, I think this is a little bit of a calming moment.

**0:34:17 AB:** In addition, there's this question around adaptation. So a lot of these estimates about how bad the future will be hinge really centrally on this question of like, "How does humanity adapt?" And so, that's the emphasis, it's like sea level rise. If you continue to build slums right on the border of the ocean and the ocean rises, a lot of people are going to be flooded. If you build berms and dams to protect the sea level then many fewer people will be hurt by sea level rise. And so, how humanity responds and how the development proceeds plays a really, really big role in estimating the impact of these kinds of climate changes. And so, we came out of this thinking like, "Climate change is really bad. We should do something about", it but it doesn't seem sort of colloquially as disastrous as some of I think, the rhetoric might imply. I'm happy to take more questions about some of these specific examples towards... After I finish.

**0:35:09 AB:** So, moving to the second question. This question like, "What could you do about it?" This is the aspect of climate change where I think we know the least. Most of the philanthropic work around climate change is some kind of advocacy spending. In the US, you're familiar with the advocacy organizations who work in sort of the Federal Climate Legislation, but there's a whole variety of different kinds of things. The Climate Works Foundation, which is one of the biggest foundations that work on this, their sort of like typical work might be campaigning for a higher efficiency standards for concrete manufacturing in China. So it's still advocacy, but it's a very different kind of beast from lobbying in Congress. And so, we don't really have a good sense of what the expected returns this kind of activity might be, and so that's, I think, a big opening and one where we'd like to do more research in the future, to try and look back at the evidence to say, "We spent a bunch of money on advocacy over the years, how successful has it been? What have we really gotten globally?"

**0:36:02 AB:** Obviously there's some questions about sort of success in the US. And then, there's another sort of approach or intervention where we think the climate community, which is fairly large, might be not paying enough attention, which is this thing called "geoengineering". It's basically a set of strategies designed to try and deal with some of the worst impacts of climate change, but not necessarily involving cutting emissions right now. And so, we think there's a potential need for more research on these issues and it's something that there might be sort of a structural problem where the climate change community isn't paying adequate attention to the status quo. Finally, the sort of third issue I mentioned is this question of, "Who's currently working on the problem and how much money is being spent on it? What does the sort of market for donations and the status quo look like?"

**0:36:43 AB:** And, this is tough to really get a sense of. Climate change involves multiple actors, there's also a political issue and so, in addition to looking at philanthropic spending, you also want to consider sort of political donations, and that's really difficult to do. Our best estimate is that around \$500 million a year is spent in the US on sort of combating climate change in various ways. Often from foundations, as I mentioned, Climate Works, they're an especially large player in this field relative to most other fields, as in they represent more of the field than any 100 typically represents in other fields. And then, the US government also spends around \$8 billion a year on climate change, about three quarters of that goes to clean tech research, basically. So, we don't have a really good sense of how this sort of \$500 million a year figure from philanthropy compares to other causes because it's not in compatible terms. It's not really something that's easy to compare,

but it seems like it might be... If you compare it to migration, which is something else that we think is a really big issue, you're talking about at least an order of magnitude, more funding for climate change than sort of different kinds of migration advocacy.

**0:37:47 AB:** There's still a lot more work to do on basically all of these questions and we think we have more to learn but this is sort of a good example of a fairly... This is an area where we know more than most of the other shallows that I mentioned on this sheet and something where we feel like we have a couple of specific questions we want to move forward with, but where we've learned a lot in our investigation.

**0:38:06 EH:** So, just two quick things before we go to questions, which is one, Alexander gave a summary of what we've learned, but our full write-up on all of this is on our website with the footnotes and the sources, and all of the data, so you can check that out. And then the other thing is, we're happy to take questions on climate change or any of the other things that are on this page that we've looked into. We present climate change not as the most important one, but just as one of the things that we had looked at, that's a good example of our process.

**0:38:33 S?:** So, as part of this investigation, have you made any progress or learned anything about sort of the long tail risks of climate change? So, the idea that there's sort of an immediate impact and then there are unlikely but very bad events that could happen, and I've always found it very hard to assess how likely those are.

**0:38:50 AB:** Yeah, so I mean we've talked to a few different climate experts about this topic and it's actually one of the sort of like sub-questions that this summer, we have some more staff capacity, we're going to finish investigating, 'cause... Well, finish, but learn a little more about and write up, because it is, I think, a lot of smart people who study climate change think that if you're sort of worried about the damage from climate change, those sort of like large tails of uncertainty, actually sort of dominate the expected value calculations. Like, how bad do you think the 95th percentile chance might be, could totally drive your estimate of how bad climate change will be. And it's really hard to get a good estimate of what the 95th percentile is.

**0:39:25 AB:** And so, this is something that we're worried about and thinking about. It's one of the reasons we're interested in geoengineering research, is because some of those tools might be able to cut off some of those tail risks. And so, there's a lot of questions around this. It's definitely, I mean, we're excited to learn more about. I don't think you're going to get... Any philanthropy scholars will tell you they're not going to be able to give you sort of good estimates of what those risks are actually likely to be. And so, globally from a humanity perspective, what you want to do is ensure against that. You want to like not have so much climate change. So, that's kind of where we are.

**0:39:55 S?:** I think the one thing that I think it might be possible to get with more investigation is a better idea of sort of the strongest arguments against and in favor of the tail risks being large or small, 'cause that beyond sort of general intuition about, you push a little in your system a certain amount and does it break or does it stay in the linear range? I don't really have any good sense of this.

**0:40:16 AB:** Yeah, I mean, I think like in areas where there's not much scientific consensus and that describes some of these areas I think, it's going to be really hard to figure out. Especially if you're us, how to moderate those sorts of disputes like, yeah, you could like... I mean, like, for instance,

there's very specific disputes about like the physics of ice sheets in Antarctica, and what kinds of conditions under which they might collapse? So, in the write-up, we sort of bounded it according to some expert claim about what the reasonable consensus was, but there are definitely credible experts who disagree with that sort of claim about the reasonable consensus is. And so, you're sort of placed in this position where you've got some radical uncertainty, you don't want to think it's infinitely valuable to rent that sort of slim chance and so you need to make decisions and it's tough.

**0:41:01 S?:** Yep. No, it's definitely very hard.

**0:41:05 S?:** So you mentioned the discount rate before we got to another topic and it seems like climate change, is if this happens, it's fate, because most of the projections are out to let's say about 50 or 1,000, 100 whenever, and one can assess the damage that's happening then, if one continues those projections to 2250, for instance. I mean, those 95% tail risks now become sort of expected.

**0:41:34 AB:** Yeah, so...

**0:41:35 S?:** 2250, obviously, that's 250 years away.

**0:41:38 AB:** Yep.

**0:41:38 S?:** So, how do you think about that?

**0:41:40 AB:** Yeah, so the question is sort of like, how do you think about discount rates and time horizons especially in climate change where the worst impacts of climate change are going to be in the fairly far future. Maybe, like, after 2100 or something like that. And I think it's a tough question, like, sort of... What we actually did in this report where we looked at the impacts was to go out as far as the impact projections went, so we took the furthest projection that was available in the IPCC report, which in most cases was around 2,100. So actually if you look at it through big economic analyses of the impacts of climate change which actually is that the worst impacts, the ones that contributed most to how much we should pay right now to avoid climate change, come in like often after 2200, like after the year 2200.

**0:42:23 AB:** So you really have to think like, "Man, we're having a lot of benefits for that generation after 2200 to think it's worth paying a lot right now to grant that." We don't have a 3% discount rate to give you. We are trying to use the best data available on the humanitarian impacts in the far future, but it's whole... It's written in the blog recently. We're fairly optimistic about what the future is going to look like in terms of economic growth and other kinds of social progress, and so I think underestimating how good the future's going to be might be due to under-investing in solving today's problems, per se, is poor. So, we're trying to balance those and it's not an easy sort of 3% answer, unfortunately.

**0:43:03 EH:** Another thing I'd want to add in there, and this is something that's going to be maybe our recurring theme, is that we don't want giving recommendations that change completely when you change the discount rate, that's like exactly what we're not going for. We're very quantitative people, but I think when... We learned a lot about what happens when you over-quantify or when you put too much weight on one kind of master quantification. And so the question is, "Can you be systematic, can you be transparent, can you be reflective without doing everything in terms of one expected value equation?" And I think the answer is "yes". I mean and we're going to be writing

more about it.

**0:43:38 EH:** One analogy is there's a big difference between, let's say you're trying to buy a printer. There's a big difference between seeing an ad and buying that printer because it looked cool, versus doing some sort of systematic search, explaining what you did. Let's say, you're doing it for someone else, you're an assistant and you look at a lot of options, you list the different properties, you separate them in a way that's easy to read, and in the end you make a judgment call that's much more informed, and it's that second version that we're going for. And the third version would be just like creating one printer goodness metric, to try and making it and just publishing the numbers. That is also not what we're going for.

**0:44:11 AB:** I want to actually add just one more detail about this. So, a lot of the disputes around the right discount rates using climate policy, the sort of philosophical justification for using discount rates in sort of social welfare analysis. Typically, I mean there's obviously the time value of money but the reason we would actually normally give is that we expect the people in the future to be much, much better off. And so if we're talking in dollar denominated terms, we want something to account for that fact. 'Cause if everybody's in the future's going to be millionaires, we don't want to pay them \$100 now for them to have a \$100 more in the future, 'cause of the declining marginal value of money. It's just not worth as much in the future if you already are all millionaires. And that's the main justification that people give for this. I think it's just interesting that like, if you actually look at the sort of economic analyses, like people in the future do look really rich, even in the developing world, in a lot of cases. So there's a lot of opportunities for growth. And so I do think having a little bit of conservatism and saying like, "Yeah, maybe we do want low discount rates because we do really value that welfare of the people who are going to be much better off, and that we should give a lot more to people who are poor right now," seems like it make some sense. But...

**0:45:16 S?:** So, you mentioned that in some cases, the development... The economic development mitigates some of these issues. But in some cases some of the interventions to help reduce climate change inhibit that very economic development. So are you looking at the sort of models that examine the direct tensions there?

**0:45:35 AB:** Yeah, well I mean I think...

**0:45:36 EH:** The question...

**0:45:37 AB:** Sorry, yeah, the question is, are there cases or have we looked at the cases where interventions around climate might limit development? And I think people talk about this in both directions. A lot of people worry about the sort of growth of low income countries leading to more climate change. So, this isn't something we have super actively explored. Obviously, it'd be a huge issue if you were thinking about which particular climate interventions you're wanting to fund. Elie published... We heard some story recently about the policies of the... Do you just want to just tell the story? The cat and mouse story?

**0:46:10 EH:** Yeah, I mean one of the... I think this is what you're referring to but in one of the causes we've been looking into on this list is providing developing world infrastructure, like better Internet, electricity, communications infrastructure. So, someone told us a story where there are certain regulations that US companies need to meet in order to invest in... That the US Trade

Organization forces them to meet in order to invest in that country, and the one major regulation that is imposed is a limit on the carbon emissions that the investments can have. These regulations were essentially pushed into place by the environmental lobby, 'cause it's a very powerful lobby. There is no lobby that is in favor of the poor African, saying, "Even though there is this global problem of carbon emissions, there's also a local problem of lack of electricity that's causing humanitarian harm." And so the more powerful, rich country lobby beats the non-existing, poor country lobby. I mean, it's also an interesting example of the type of political dynamic that Holden was talking about before.

**0:47:20 HK:** One more quick answer is that the report we've been using, the IPCC does try to answer this question in a very general way. I mean, it gives scenarios with high growth and no attempts to mitigate climate change, and low growth and no attempts, and then some attempts. In general, I mean the kind of pattern is like, the best thing that can happen, period is just like really high growth, kind of with or without attempts of mitigation. But the way they're modeling it, the attempts of mitigation, we still come out like really good. The thing that's worse is that kind of slow growth and no attempts at mitigation. So...

**0:47:55 S?:** What do they talk about, whether the attempts of mitigation effect the likelihood of high growth versus low growth?

**0:47:59 EH:** I believe they're trying to model that, right? In those areas...

**0:48:01 AB:** Yes. And so, it's not like... In 2007, that wasn't so sophisticated. I think people think like, there's nothing like a necessary trade-off. Like the cost of using, trying to use more efficient technology is relatively cheap in the scheme of things. You're not committed if you use clean tech to slow growth, but...

**0:48:19 S?:** In the GiveWell Labs projects', in the research, how much are you kind of prioritizing research that you're trying to understand like the whole macro picture, all the macro trends and research, into just finding kind of weird opportunistic, giving opportunities that just are low hanging fruit and make sense, that maybe you under... I ask 'cause it seems so hard to quantify, maybe not even necessarily quantify, but understand all the macro issues and then make a decision, whereas this may be some low-hanging fruit that you don't quite understand the macro trend as well, but you know that there's a lot of aspects of this given opportunity that make it a good one that you will invest. So, how do you...

**0:49:06 AB:** So, I mean most of the work that we're... Sorry. [laughter] Repeating the question. The question is, how do we sort of think about these big picture investigations, versus sort of lower depth investigations where we might already know more about something or not need to know as much, to find a good giving opportunity without having fully investigated some giant cause in this broad stroke terms? Is that a fair... Is that right? Yeah. So, I think it's a tough one. Most of the work we're doing is these sort of broad stroke investigations and it's for the reason Holden sort of laid out, about needing to sort of feel like you understand the cause, to really believe that any particular giving opportunity is a good one. And that said, we actually do have some sort of theories that are a little weird or that might be sort of idiosyncratic, that we think might be... You need to account for why something is overlooked and why it might be a good giving opportunity, even in some kind of field.



**0:49:54 AB:** And so geoengineering is the best example, where even if we're not so sure that marginal climate change philanthropy is where you want to invest, geoengineering might be useful because climate change philanthropists come from this environmental background and don't want to mess with the earth, which is... It's not an unreasonable perspective, but it might lead them to systematically underinvest in the kinds of research around second best opportunities to mess with the earth, through that sort of really catastrophic climate change. And so that's an example of where, if you have a theory about why something wouldn't be working in its field, then you might want to look at it, like a sort of specific opportunity more closely.

**0:50:28 HK:** I think the answer is sort of both in other words. It's that, if something important enough, then we'll look at it, regardless of whether we know of any particular missing opportunity, because we may... I'm looking into it, finding those missing opportunities and if it's important enough, you want to invest in that. And the other way a cause can get on our priority list is if it's just a really strong case that there's something that's being ignored. So, it's really either one, but I think by... We'll definitely want to look into anything that is claimed by many people to be very important and that, I think, over time, we're going to build up a pretty decent picture of what can be known about the macro trends, which is very well...

**0:51:03 EH:** But a lot of the work we're doing right now on these issues, is in itself is experimental, because we're trying to figure out, how much do we have to look into something to move far enough in our understanding that we can better prioritize it against other things that we could look into? So, Alexander is talking about climate change. He spent time on it, but in the scheme of things, we know very little about climate change. And the question is, how much do we have to know about it and the other eight things on the list and the other 50 things we're going to research in the next year, in order to make a better decision, a decision we're more comfortable with, about choosing the... I don't know... The three areas, the four areas, that we're going to put more of our time into. And that, I don't know the answer, I don't know whether people have done that before. If they have, they don't... It's not shared publicly, so we feel like we're trying to figure that out as we go on.

**0:51:52 S3:** And regarding individual projects and people who seem really promising and overlooked, there's just... There's often reasons... There are a lot of philanthropists that are looking for unorthodox approaches and outstanding social entrepreneurs and trying to pick that low-hanging fruit. And oftentimes, if there's a sort of... And something that seems to be a gap in an area, the philanthropists who's focused on the area may know about it, but may not be filling it on purpose. And so, it just seems like you need, again, a lot of context to be able to say, "This is genuinely overlooked for bad reasons or not." Or... Yeah.

**0:52:41 HK:** It's really easy to be dumb money.

[chuckle]

**0:52:47 S?:** Just quickly on climate change, how do you treat the fact that it's sort of often thought of as like an all or nothing solution, so like having a little impact on climate change doesn't get you very far, while having a little impact on some other stuff, you can save a lot of people's lives?

**0:53:07 AB:** Yeah. The question is, how do we treat the fact that climate change is often thought of as all or nothing? Like maybe you get to two degrees or you don't, and it's not so much like any

marginal spending is going to get you some marginal benefit. So, my impression is, I've talked to a couple of sort of climate funders on that and it's like... My impression is that it's basically wrong, that there's this worry about climate tipping points, but actually the more climate change there is, the worst impact it has likely to be. And so, getting from 2 degrees to 1.5 degrees is actually not as good as getting from 2.5 to 2 degrees. You really want to prevent the marginal changes whenever you can. So, this idea that there certainly are unpredictable non-linearities, don't get me wrong and so you want to avoid going over those, but because they're unpredictable, we don't know exactly where they are. There's not some sort of bright line where it's like, "Oh well, we're already going to go over it, it doesn't matter anymore." The worst the climate change is going to be, the more you should expect yourself to be going, to committing to these for a non-linear areas.

**0:54:02 EH:** Let's just pause on climate change now, so I just want to kind of offer the last section of information on what we're doing and then open it up again for questions on the things I said, or any of the other things we've discussed, or are on this page before we wrap up. So, the part that I want to talk about is what we're doing now on our... Some of the things we're doing now on our more traditional work of looking for developing world direct aid organizations, these proven cost effective opportunities. And I want to just focus in on two specific things that we've done recently that I think are somewhat interesting. So part of our work is looking for charities that are competent, and running programs and able to show that what they're doing is working. But a big part of looking at international aid is not only looking at the charity, but also looking at the program that it's implementing. So you have the Against Malaria Foundation, they're an organization, they deliver bed nets to prevent malaria, and so another part of our research is understanding how good bed nets are, how much do they cost? How effective are they? When do they work? When do they not work? And what can you say about, let's say, the cost per life saved?

**0:55:10 EH:** So recently, there are two programs that we've been looking at and these are only, I only share them because they're interesting, not because they're a particularly big part of what we're doing but I think they're an interesting example of how we work. So one of them is Vitamin A supplementation. There's this program where if you... There's a good number of randomized control trials showing that if you deliver a vitamin A supplement to children between the ages of six months and five years, twice a year, you reduce child mortality. There's a number of studies that show that child mortality falls by 25%. The supplements are incredibly cheap, so this looks like a very good program. There recently was completed a very large trial where they delivered vitamin A supplements. They had a treatment group, a control group, and they found no effect on vitamin A supplementation at all. So, if you...

**0:56:02 S?:** Child mortality, you mean?

**0:56:03 EH:** On yes, on child mortality. Sorry. If you... If you take all the results together and you just mathematically average them, you'd go from maybe 25% effect to a 12% effect, but it's not quite clear that that's the right step to take with all of this data. The initial vitamin A trials are all from a decade or more ago. This more recent trial obviously happened more recently and over the period of time between the initial trials and this larger trial, child mortality fell a very large amount throughout the developing world. I mean, you have this overall trend but measles deaths, in particular, go from about a million a year to 200,000 deaths a year, over the period of, of different time periods between these two trials, and so there's some question of whether the opportunities to provide vitamin A supplements and prevent child mortality still exists today. So, this is actually something we're literally in the midst of right now and we don't, I don't have a conclusion to offer

you other than to say, "This is the type of work that we're doing where we're looking at the evidence and trying to sort through these thorny questions of, what does it all mean? The actual next step in our process is talking to some of the researchers who have worked on this, who have their own views about what the implications are of this more recent trial.

**0:57:27 EH:** The other area that we've spent some time on this year that's interesting is the program of water purification. So when you hear about developing world water causes, some people are talking about digging wells, other people are talking about providing sanitation infrastructure. This program is purifying water. It's putting chlorine into water to reduce contaminants or it's creating filtration systems that make the water purer. And the intention of these programs is to purify the water such that it reduces contaminants and prevents childhood diarrhea, which is one of the leading causes of child deaths in the developing world. So we looked at this evidence and again, there's been a large number of randomized controlled trials and collectively, they show a pretty sizable impact on childhood diarrhea rates. But the case is not that simple because one methodological distinction between the trials is some are blinded, meaning that the people who are treated know that they are the ones who have received the cleaner water and the ones who haven't, do not, and they're reporting diarrhea based on self reports, it's not objectively assessed data.

**0:58:42 EH:** And there's a fear that people are answering the researchers in the way that the researchers want to hear. So the people who got the chlorinated water say, "Yes, sure, it's had a big impact", and the people who haven't are less likely to say it, and this concern is borne out by the fact that the blinded trials, where they maybe gave people a filter for their water that doesn't actually have a purifying agent, so it shouldn't clean the water, those trials, there's fewer of them, but they don't find an effect on diarrhea rates. Of course, even that is not that simple because these... I don't think it would be fair to say that we know that this program does not have an effect because in those three trials, there are some problems. One of the problems is that the control filter, they measured what it did to the water and actually purified the water, and so there's some problem that it...

[laughter]

**0:59:37 EH:** It wasn't quite as bad as they thought it would be. There's also an interesting argument and it's a paper that we actually just came across in talking about what we had found with someone who is very familiar with the water literature, that just by asking people about whether their child has had diarrhea in the last week and asking them repeatedly over the course of a trial leads to them changing their behavior, in a way that prevents their children from actually coming down with diarrhea. And so one of the things that you actually see or to more expensive intervention to send someone to ping you every week about what's going on, but it's actually interesting. If you look at the... One of the more recent blinded trials, you can look at a chart of the control group and the treatment group, and reported diarrhea, and both of them just fall linearly over the course of the trial. It doesn't really make any sense that that is how it works, that full of treatment and control would start high and slowly move down except perhaps for the fact that another intervention is the one that's more effective at preventing this problem.

**1:00:42 EH:** So, again this is something where our research is further along. We have a draft page put together but we only recently share what we had found with someone who is very familiar with water and they gave us some more research to look at, and so we're still in the midst of finishing

that up and then we're going to publish that all on our website. And so now, I just want to open it up for questions on what I said or anything else. Yeah, Darren?

**1:01:05 S?:** For the vitamin A studies, if you look across all the different studies that disagree with one another, do they have any differences in dosage? 'Cause one thing I've read a little bit about is that at very high dosages, vitamin A has sometimes been associated with increases in mortality, at low doses, it's always been associated with decreases. Most of these are developed world studies, but I wonder if that effect might be a plager as well?

**1:01:27 EH:** Yes, so the question is about whether there are differences in the size of the doses between the different sites. I don't know the answer to that question. Tim, do you?

**1:01:33 Speaker 5:** Yeah, actually in the vast majority of the trials that were done, they were all based off of world global guidelines and so they were actually of varied resistance but hasn't been tested very well.

**1:01:46 S?:** But all the studies were something like 100% of RDA, then actually in this case that would be that this wouldn't be a concern for these studies?

**1:01:54 S5:** Right.

**1:01:56 HK:** The big worry there is that the big study that Elie was talking about, like people think it might have been really poorly conducted.

**1:02:03 EH:** Yeah, these are things that obviously, like, are there differences in the trials that could explain the different results? The major thing people point to is that this large study had some issues but I don't really know what those are yet. We're looking into it.

**1:02:15 S?:** So, you get access to the raw data?

**1:02:18 EH:** No, not yet. But I mean, this is, I mean, basically the... You know there's different things. If you go on our website you see the bed nets report, the cash transfers report, those are completed, went through the full process. The water, I would say is something like 75% done where we did a lot of work, found a lot of research, formed some view about the major questions, and then got to the point where we said... You know we had talked to water researchers early on but then also into another water researcher later on and said, "Here's our conclusion. What do you disagree with?" We want to know before we publish this. Vitamin A is like the most initial stage. We've probably put a week into it so far and see this is the major question we need to look into.

**1:02:55 HK:** We also often don't... We'll get raw data for a study and re-analyze it, if the study is that important to what we're doing, but I mean that's a huge endeavor to get the raw data, just to get it a lot of the time.

**1:03:08 S?:** Do you consider, whether or not the problem should be solved by philanthropy, by other funding, types of funding, like government?

**1:03:18 EH:** Sure. So the question is, when we're looking for giving opportunities, do we think about whether it's properly the role of philanthropy or government? I would say we...

**1:03:27 S?:** Or industry.

**1:03:28 EH:** Or industry, right. So, we certainly have those discussions informally. I mean, I think a lot of what our kind of general formula is, is looking for things that are really good and really ought to happen and really aren't going to happen if we don't step in. And it's that last criteria that I think somewhat takes care of a lot of this and it's definitely good to get into the theoretical issues, like, why aren't bed nets just sold and why aren't they just done through the government? I think we have good answers to both those questions if you want to discuss them, but the bottom line is that bed nets save lives and you can look at the data that, that coverage for years, and we're not reaching over and we're not even close. And so that, to me, is a little bit less of a shaky foundation than having theoretical arguments about these things because it's like, okay, well, for whatever reason this isn't happening. Industry is not doing anything, government's not doing anything. If we do it, people's lives will be saved.

**1:04:21 EH:** Right. So the questions we'll ask are, "If we didn't direct the funding here, would it be provided from somewhere else?" You obviously want to be somewhere where as a philanthropist, you are making a difference and causing something to occur that otherwise would not have. And then we'll also ask the question, "What are the possible negative or offsetting impacts of coming in and providing funding here?" So, every page where we write about our program or a charity, we'll have a section where we are trying to analyze those types of potential downsides.

**1:04:49 S?:** So, perhaps a related question is, who do you see... So, kudos on Labs, I mean I think the Labs expansion, the rationale for it sounds perfectly right. So, once Labs is all kind of really going, how do you see it sitting alongside your other recommendations? Are you going to rank Labs-y things relative to other things or are there going to be two different rankings, the one you have now in a different lab? Or the Labs, it's not even going to be ranked?

**1:05:18 EH:** Right.

**1:05:18 S?:** And do you see... And this goes more to her question... But do you see the audience of the Labs output being slightly different than the main reading audience of the current stuff?

**1:05:28 EH:** So the question is, how does Labs and the recommendations we make there interact with our existing current global health top charities? I don't think it's something that we know the answer to yet, 'cause we haven't gone through the process of having that, but I suspect like a few things will be true. I mean, one, I think that initially, the main audience for Labs will be very large philanthropists. Good Ventures is sort of the initial customer, so to speak, for those recommendations. But I mean, we and they both want other big philanthropists to come in and support those projects, and I suspect that a large portion of the individual donors who have made up our audience for years, will want to participate in those projects as well. I suspect we won't do some sort of forced ranking of everything, though, like one thing we'll always do is at the end of the year, I'll give to charity, other staff will give to charity, and we're going to say where we gave, and so, in some sense we'll always be forced to do some sort of personal forced ranking of where we think additional money goes the furthest. Yeah, John?

**1:06:39 S?:** With respect to policy advocacy, how big a difference do you see between that and something like global health in terms of your ability to translate dollars put into the cause into an

outcome? So, for AMF, you can say, "I spent X many dollars, I get Y many bed nets, I expect this outcome." Whereas, for policy, maybe you're talking about a number of petition, names you got on a petition, and there's also lot of... Potentially, a lot of horse trading behind the scenes where you don't know what had to happen for you to get your potential legislative victory.

**1:07:19 HK:** Right. So, the question is, how do we fit... When we talk about bed nets, it's very easy to say, "My \$5 bought a bed net, this many bed nets saves a life", and do all that. And with policy advocacy, or actually I would add with scientific research, it's going to be much harder to draw those kind of connections. So, my answer is, it's going to be a lot harder and less satisfying. And that, I think, actually is a reason that I'm kind of excited about these two things [laughter], that a lot of philanthropists are excited about them because the stuff, the more kind of emotionally satisfying and tangible something is, the more, my guess, is that, the kind of people that we're not, are funding it, and the kind of people that, the only kind of people who fund this stuff that's much harder to get your head around, are going to be our audience, who I think is generally going to be up for that stuff.

**1:08:10 HK:** So, I actually... I don't really think it's a problem. Obviously, in a perfect world, I would love to know exactly what came out of every dollar, but in the real world, where some things are much more noble than others, there's arguments both ways. There's arguments that you learn better and you can keep yourself on as better when you're measuring better, and there's arguments that when you don't have that satisfaction, that's where you really need the most sophisticated donors to step in. So, I think it's going both ways. I think we're going to try to express what we think you're buying, in the way that we think is most helpful, and that's the way that we've always done it. We've never really gone in from putting out numbers that we think are just so kind of made up they're silly. We tell you what numbers we think are useful.

**1:08:51 HK:** And so, with policy advocacy, I could imagine saying something like, "We are recommending this large amount to facilitate a substantial increase in capacity of this large organization. The reason is that, this issue does not have the kind of capacity that these other issues have." And when you look at these other issues, you see much more powerful lobbies. Now, what does all this translate to? Well, in general, we think that philanthropy has such and such track record of winning these big fights and if we want to fight this big, then this is what we would need. Well, that's a lot of steps. That's not the world's most satisfying thing but that's to me like, what form an actually reasonably good answer might take.

**1:09:34 EH:** Yeah. I mean, there's one thing in there I would say a little bit differently which is that, one of the... Holden said earlier, one of the reasons we started with the global health measurable stuff is that you would know at the end of the day whether you were right about what you thought or you were wrong, and you could be held accountable. And I think it's something I feel like I'm personally focused on or concerned about, that it can become easy to get lazier when you don't have that same type of bottom line accountability for the recommendations that you make. I think that's something that is definite, I guess, something we're aware of and it's just something that we're going to have to continue to be bringing the same level of rigor and critical thought to less quantifiable areas, as we have to the more quantifiable areas. And that's something we'll figure out but I hope, I don't know, like you hold us accountable to the fact that we're doing that, 'cause we're going to keep publishing what we think and what we're recommending and why.

**1:10:30 HK:** When we get to these long time horizons, I think debate and transparency become

even more important because we no longer can follow a charity and be like, "What happened?" The best form of accountability we have is people arguing with those and pushing back to us.

**1:10:45 S?:** Okay. So, it's really tied in with the question of making it difficult to quantify things like political advocacy, it seems intuitively, it makes sense that an organization of all... Organizations that have more money is more likely to be corrupted than an organization with less money. But there's probably also factors that particular organizations do, in that they are more effective at doing things, even when they're given a quantity of resources and giving a set of people an amount of money. Is that something that you plan to study, and if so, what would you do with that information?

**1:11:14 HK:** So, the kind of observation is that, maybe all else is equal, a better funded organization is more effective but there's also a lot of factors that make some organizations more effective than others that don't have any... That are independent of how well-funded they are and how we're going to consider this. I think the answer is that, this is one of the many, very tough things about philanthropy that makes it very hard to capture in a kind of satisfying equation or system, is that, you could never fully separate... And anyone who does I think, is making a big mistake. You can never fully separate the intervention, the thing you're trying to do from the people and the organization and the culture of the organization that are going to be doing it. And so, you have to make judgments on both, you're often making judgments on both at the same time and they just both have to be factors. And I don't have much to say about it except there's not much choice.

**1:12:02 HK:** So you can see a preview of how we handled it with last year's top charity announcement because we had... We had a situation where if you look in isolation at deworming versus bed nets versus cash, I think I'd probably rank them... According to our cost effectiveness analysis, number one, bed nets, number two, deworming, number three, cash, and bed nets and deworming are closer. But there was this fact that the bed nets and the cash organizations were really transparent, really communicative, just superstars in terms our ability to be confident in how the organizations were running, how they were operating, knowing we understood them, knowing we're able to communicate with them.

**1:12:43 HK:** And while I think that the other one is an outstanding organization in this scheme of things, it was not at that level and so we didn't rank them that way, we ranked cash, number two. That's our preview and we explained why we were doing it and we listed the observations that led us to think they were better organizations. Now, I don't think there's a very robust academic literature that will tell you like, "When you see X in organization that means it's good and when you see Y that means it's bad." I'm pretty familiar with social science, I don't think that's going to happen, but I do think that we are going to be definitely thinking about this, and always spelling out our views and when we can take advantage of empirical data on whether it's a good organization, then we will. And when we can't, then we'll use the expert opinion, we'll do that, too.

**1:13:23 S?:** It seems a little bit analogous to the meta research questions, it's kind of meta-policy. And like, how do you actually think an effective organization can do this?

**1:13:29 HK:** Yeah and we're always running all the questions into each other at the same time. A really good giving opportunity means that you've found an important problem but also something you could do about it and also, it's not already being done and it won't be done if you don't it, and the people you're fronting to do it are up to the challenge. That's a lot to ask for any giving

opportunity and that's why we intended to go pretty deep on investigation before we recommend.

**1:13:52 EH:** Let's take one more question and then we'll... Let's hear more questions?

**1:13:55 S?:** The meta question. What about yourself? Right? I mean as Labs expands the set of things that might be good for the world, so I claim GiveWell, in this different world, I claim that not, like you said, in the beginning, Holden, that not everybody thinks this way and a thing that would be good for the world that might be on the Labs list is to get more people think this way.

**1:14:18 HK:** Right.

**1:14:18 S?:** So, is GiveWell going to rank as a cause, getting more people to pay attention to GiveWell or scaling GiveWell or replicating GiveWell?

**1:14:29 HK:** So number four...

[laughter]

**1:14:31 HK:** So, the question is...

**1:14:32 S?:** Great question.

[laughter]

**1:14:34 HK:** Yeah. The question is, where and how do we rank GiveWell next to our other giving opportunities if we think that what we're doing is so great? [chuckle] That might be one definite way of putting it. I think that we do have to take a fundamentally different approach to asking for money for GiveWell itself because there's actual conflict of interest issues and I think that what we have done historically is what we're going to continue to do which is we... In any kind of public outreach manner, we're never asking for money and we want to always avoid asking money.

**1:15:04 S?:** So to be clear, I wasn't so much talking about about your guys trying to get more money, as much as, are you trying to scale your own organization or maybe make copies of it? Start embryony or throw out seeds like a tree would and get... I don't know some copy on other countries to do, to some organizations, or just grow?

**1:15:23 EH:** I think like our vision, long term is that, there's way more people following GiveWell and giving to the types of things we recommend. There's other GiveWell for other causes that we never look into. But I mean frankly, we have bitten off in some ways more than we can chew just with the research we're doing, and so our focus now is like, "Do this research well. Show or model the way to do this research beyond just the quantifiable measurable stuff", and we're very... That I think is very much in an experimental phase right now and if we can do that, I hope that we'll have the capacity to go into these other areas. So we are growing GiveWell, and we're small team. We could potentially double in size by the end of the year with people who are sort of in consideration for being hired. We'll see were we end up but our major focus now is growing that research capacity, so we can do more of this analysis that we...

**1:16:16 S?:** The similar part of it, is there some way that may get more people to think this way, in



general, without just growing an organization? I don't know... Did you hear about the Wall Street Journal article on the Arnold Foundation?

**1:16:28 HK:** Right. The question is, is there a way to get more people to think this way that doesn't involved just growing GiveWell? I would basically say that's not where our heads are at right now, just to say that we're trying to focus on what we do best and that's largely research. I think that other people have done a better job marketing us than we can do for ourselves. If we ran into someone, I mean, as a giving opportunity, if we ran into someone who's really good at promoting this way of thinking, we thought they're going to have a huge impact, they needed money and they're good at it, and all that stuff that we're looking for, that would be one thing. But in terms of our mission, I mean, it's something we debate with our Board constantly and you can listen to many recordings of it, and I think it's... I think the right path for us right now, I think we seem to have a lot of work ahead of us on research to figure out kind of scientific research policies, the hundred causes we're looking at, it's going to be at least a couple weeks, so...

[laughter]

**1:17:18 S3:** I can add if you don't mind. But I think there is like this emerging movement around effective giving and I think GiveWell's existence has helped to allow that to happen, because people have a resource that they can turn to when they want to give effectively, but don't have all the time in the world to devote to research in these questions. And so I think that movement is emerging. I think without GiveWell, it would emerge a lot more slowly, so someone said, that supports what their saying, that the best work they can be doing right now is the research.

**1:17:50 EH:** And we participate in that movement, too. I mean, we hope that along, when there's kind of high return, low cost ways to do so, like speaking in an event for people in that movement. Not just this one but like, for other organizations, too.

**1:18:01 S?:** Yeah.

**1:18:02 EH:** Alright, so let's wrap up. I mean, we really appreciate all of your coming. I mean, like Holden said, as we move into these new areas, the debate, the questioning, is really core to our ability to do our work, so we really appreciate your being willing to come out and, I don't know, participate in the different type of charity event than you might normally go to. That's really helpful to us. And we're happy to stick around as long as people want and just talk informally about what we're doing.