

GiveWell New York City Research Event, April 10, 2014

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.

0:00:11 Eliza Scheffler: Hi everyone. Thanks for coming. Elie is just setting up I think. We're going to record this. We try to record meetings so we can post them on our website afterwards. So you should feel free to speak freely, and if there's anything that you say that you would want excised from the recording, we can do that. So just mention it when you're speaking or come and see me afterwards.

0:00:34 Elie Hassenfeld: Alright, well, thanks everyone for coming. Thanks, Eliza, for the welcome. I appreciate that. Just before we get started, I wanted to make sure everyone knew who the GiveWell-related people here are, Eliza's a research analyst at GiveWell. She's been here at GiveWell for a couple of years. And this is Tom Rutledge, he's our board chair since a long time he's been on the GiveWell board, so a long-standing supporter. And I'm Elie Hassenfeld. I'm the co-founder and co-executive director of GiveWell. The basic plan for this evening is to have an informal conversation about the research that we're doing. The goal is to bring you all into the sort of the research room that we have where we're trying to decide what to do. So a lot of the questions that we're going to talk about are ones where we don't yet have final answers, but they're really the challenges that we're sort of currently facing as we go through our research process.

0:01:29 EH: So there's basically going to be two parts to this event. First, I'll give a brief overview of the work that we're currently doing on our traditional focus on GiveWell Top Charities, which is organizations implementing, evidence-backed programs that serve the global poor, and then we'll pause and we'll have a little bit of time for some questions. And then we'll move on to an overview of what we're working on with GiveWell Labs, which is our newer, more open-ended research process and then have some time for questions about that. And then Eliza and I, and maybe Tom, we don't know, are happy to hang around afterwards and answer any questions one-on-one, that people have. So I'll dive into sort of our traditional work time on GiveWell Top Charities.

0:02:18 EH: So this is the part of GiveWell that looks for the most quantifiable evidence-backed programs, and they're always working in the developing world. So these are programs like deworming pills to kill parasitic infections, bed nets to prevent Malaria, and this is the thing that is most in evidence on our website. The biggest change in the last year on this side of GiveWell's work is that we grew pretty substantially; our staff size did over the past year. We went from a full-time staff of five to a full-time staff of 11, and so our plans for what we're going to do on the GiveWell traditional side are significantly more aggressive in 2014 than they were in 2013. So there's basically three parts to this research. There's following up on past Top Charities to see how they're doing.

0:03:12 EH: So there are groups that have received significant funds as a result of our recommendation. There's trying to find new charities, and then there's some, I call it somewhat new experimental work we're doing to try and support the creation of new future GiveWell Top Charities. So there aren't, so far, many updates on our Top Charities from where they were at the end of last year. I think the biggest news is that the Against Malaria Foundation, which was our top recommendation for two years, they were our recommendation from the end of 2011 to the end of 2013, and we took them down because they had been unable to find a place to distribute, to spend

most of their funds, and we felt that they had sufficient funding and didn't need additional funding at the time. In the last few weeks they've committed to a distribution in the Democratic Republic of the Congo that accounts for about two million of the \$14 million they currently have.

0:04:10 EH: So this is good news. We're optimistic that AMF is moving in the right direction with finalizing distributions, but we're not yet ready to put them back on our Top Charities list 'cause we're still watching and waiting to see what they do. GiveDirectly received more than \$10 million based on GiveWell's recommendation last year. As far as we know, based on our recent conversations, they seem on track to distribute it. And there's no major update from the other two Top Charities. So I want to go over a little bit about some of the new areas we're looking for additional Top Charities in.

0:04:48 EH: One of the... So we're just going to kind of go cause-by-cause. One of the areas that we're looking in is the program called Salt Iodization. This is where people fortify salt with iodine. This is something that... The evidence states that this program has an impact, is pretty strong. If you have a severe level of iodine deficiency, there's strong evidence that reducing that deficiency leads to people not having very significant cognitive impairment. And when people have a more moderate deficiency, the fortifying their salt with iodine, leads to improvements in, more like moderate level improvements in cognitive development as measured by IQ. This is a program that was pretty much rolled out across the developed world a long time ago. There was reasonable success rolling it out in the developing world in the '80s and '90s, and some of what we've heard from the groups that we've spoken with is that now, donors are fatigued with salt iodization, and to some extent feel like the job is complete and have moved on to other causes like AIDS or malaria, and that has left this cause, which is extremely promising without sufficient funding to fully complete the job that was there.

0:06:06 EH: So this is something we're still in progress on. I wouldn't say we've fully vetted that particular story, but there are two groups that we're looking at in salt iodization. One is ICCIDD, which I think stands for the International Council for the Control of Iodine Deficiency Disorders. They're basically a network of academics, and what they do is provide... They advocate and provide technical assistance to governments in their implementation of salt iodization. So that can be convincing them that they should do this if they're not already doing so, or providing answers to questions when the government finds out that there's some problem in their iodization scheme.

0:06:48 EH: The people from ICCIDD are there to provide support. So this is definitely a different type of group than the ones that GiveWell has traditionally recommended. They're not a direct service organization, they're advocacy. But they seem to be a group that could possibly be having a very, playing a very big role in salt iodization. The other group is called GAIN; they're the Global Alliance for Improved Nutrition. They're funded heavily by the Bill and Melinda Gates Foundation, and they run a wide variety of nutrition programs. But we're looking at them specifically for iodization, and where ICCIDD focuses more on the, I'd say, high-level advocacy and research support, GAIN is more involved in a more direct, advisory level, helping governments with things like insuring that the quality assurance checks on the salt that is iodized are appropriate, and that they can determine that they've iodized salt sufficiently.

0:07:50 EH: Certainly a challenge that we'll face with GAIN, is because they're on so many programs, there's always the question of whether a donation to the program itself supports more of the program itself, or gets, is fungible with other activities of that organization. But given the state

where GiveWell is today, we think that this problem of figuring out how to work with organizations that run many programs is an important question for us to tackle.

0:08:20 EH: We're also looking at a new organization that's part of the family of causes known as "neglected tropical diseases." Two of our charities now are de-worming organizations. Those are... That's one set of neglected tropical diseases. This organization, which is known as the Center for Neglected Tropical Diseases, actually focuses on Lymphatic Filariasis, which, whose worst negative impact is a swelling of the limbs through Elephantiasis that leads to severe disability; people are basically incapable of moving around. And so, this again, in many ways, fits the bill for a potential GiveWell Top Charity, and so we're looking closely at them.

0:09:02 EH: And then finally, we're revisiting organizations that focus on immunizations. We looked at this two years ago, and a lot of what we found then was that it didn't seem that money was the core bottleneck to immunizing more people. Only about 80% of people are fully immunized worldwide, so there certainly is room to get there, but it didn't seem like it was a direct funding bottleneck. We want to revisit that conclusion again, and so we're doing so, looking at a few organizations: UNICEF's Tetanus immunization campaign, a measles and Rubella initiative, and then also this campaign to roll out the Meningitis A vaccine, which is new and was funded, the development of which was funded by the Gates Foundation.

0:09:49 EH: Another particular challenge with immunizations is that a lot of the cost-effectiveness of the program is heavily dependent on local conditions. And that's information we would need from the charity itself. So just to give an example quickly, the way the Maternal and Neonatal Tetanus Initiative works is that they provide a supplemental immunization to all women of childbearing age. And the problem is that this... The need for this program varies widely depending on whether or not those women have already been immunized. If they have been, they're not getting a very much benefit. If they haven't been, they're going to get a great deal of benefit. And routine immunization rates vary very widely across the world.

0:10:31 EH: Similarly, Tetanus is something that is a worse problem in hospitals that are unclean and less of a problem in hospitals that are very clean, and so the likely impact of reducing Tetanus, or reducing the possibility of transmitting Tetanus, is going to depend on the levels of tetanus in that community, excluding the immunization. So this will be a challenge of looking at each of these entities if we get to that point.

0:10:57 EH: The final area of our work is this, what I refer to as a more experimental part of our process, which is trying to support the creation of new GiveWell Top Charities. And so, we've done a couple of things there so far. First, is we've... When I say "we," I often mean GiveWell and working in consort with Good Ventures, which is a foundation we work very closely with. Good Ventures provided a grant to an organization called New Incentives, and New Incentives runs a conditional cash transfer program in Nigeria with the goal of reducing mother to child transmission of HIV. If mothers meet the conditions, then they receive a cash transfer. This organization is extremely young; they're like a year old. They have basically no track record to speak of, but they are implementing a priority program. They seem transparent to us and so we're experimenting with the idea of providing them some up-front support as with the goal of having a new top charity five years from now.

0:12:01 EH: The other side of this... So one side of it is organizations implementing priority

programs, the other side is the evidence base, what programs have sufficient evidence behind them? And so the other thing that we've been trying to do is support the groups that run rigorous trials of development programs to see if we can add to the evidence base that certain programs are effective. And so we've talked to Innovations for Poverty Action, which is the research group at Yale that runs the randomized control trials. And we've also talked to a younger, smaller group that instead of focusing on, I would say, purely academic research within development focus, they work more closely with organizations on the ground to design evaluations of their programs that can... To evaluate whether or not they're working.

0:12:48 EH: And so, we're hopeful that through this process, and I'm happy to sort of go into this more, and the details of the programs that we're most interested in learning more about, we're hoping that this leads to additional GiveWell Top Charities down the line, because certainly something that we feel we've run into over the last few years, is a lack of organizations that are running the programs that we think have the strongest evidence behind them.

0:13:13 EH: So I want to pause there. That's kind of an overview of what's currently going on the traditional side of our research, and answer any questions that you have about that. Yep.

0:13:26 Speaker 3: How are you looking to evaluate, for instance the charity... The one with the iodine, the research and support side. What sort of metrics are you looking to evaluate whether that... On the cost-effectiveness side to that?

0:13:40 EH: Yeah, so the main things that we're looking at there... So our research process is basically divided into two parts: There's research on the intervention, which is, how good would it be if we accomplished salt iodization? What do we generally know about the costs? What can we say about how often these programs work? And that is totally external to the organization itself. And so we've done that for all of our top charities that exists. And then the other side is assessing the organization itself. And I think in this case, a lot of what we'll be looking at is, what can we learn about their track record? And by track record, I mean like our best understanding of the significance of the role that they played in salt iodization historically.

0:14:30 EH: We'll also try and talk to other people who are not them, to triangulate what role they played. I don't think we're going to be able to get to the point of turning... Having a very credible estimate of their dollars in and what comes out, because clearly they're one tiny part of the process. And I think the way that we'll end up thinking about cost-effectiveness is what are the total expenditures that are necessary for implementing salt iodization, and what is the total benefit, and then trying to determine whether or not we're convinced that they played a significant role.

0:15:06 EH: Thanks.

0:15:09 EH: Yep?

0:15:10 S3: This model seems very similar to the ones used by Deworm the World, and it seems like it has the same sort of challenges in terms of figuring out what the cost effectiveness is. Do you feel that the role of ICCIDD or GAIN, is similar to DTW, and if the challenges are also the same?

0:15:29 EH: Yeah, so I'm going to repeat the question just for the recording. I skipped this one, 'cause you were close. So the question is, how does ICCIDD compare to Deworm the World, or

ICCIDD and GAIN, compare to Deworm the World, which is a charity we've already recommended? I think in many ways, it's similar, in that they're not the direct service provider; they're an advocate. But I am... I don't think I'm... I'm not confident that the type of advocacy they provide is at all similar.

0:15:59 EH: So Deworm the World is... We have a decent understanding of the role that they're playing. They're very involved in the Indian states where they work, and provide a great deal of support. Often, or even providing funding for essential parts of the program, like monitoring or training or surveying to determine where deworming is needed, and ICCIDD by contrast, is much further away from what's going on on the ground. They're a network of academics; I don't think any of them are full-time employed by ICCIDD. Their total annual budget, up through this year was about a million dollars a year. They think they... They lost some funding because of some changes in Australia, and they think they can expand, but they are acting at a much higher level in some sense in the process than Deworm the World is, and so that will be its own new challenge for us.

0:16:59 EH: We've talked least to GAIN so far, so it's difficult to say exactly what role they play. My instinct so far is that it's closer to what Deworm the World does, where the big challenge with GAIN will really be this issue that they run so many programs, and we'll really have to be sure that funding coming into their iodization program causes that program to expand as opposed to something else. Yep?

0:17:30 Speaker 4: Are you concerned about... Since you'll be funding future top charities, are you concerned about your ability to remain unbiased and your ability to judge operative organizations or you might have personal stake in [0:17:45] .

0:17:46 EH: Yeah, so the question is whether we think we'll stay unbiased when organizations we supported from the beginning are ones that we're considering for recommendation. I think it's a good question. It's not something that I'm terribly concerned about, nor is it something you should be terribly concerned about. I think we've a pretty good track record of critically assessing the judgments that we previously made and being open to saying where things went wrong. You see that with what happened with... I mean all the way back, PSI and VillageReach and AMF and SCI, and so we've a good track record of that.

0:18:24 EH: There's also something that's really built into the DNA of GiveWell which is, critically assessing the information that we have and being open about it, as opposed to trying to confirm that what we previously thought is correct, or protect our position as someone who, as a group, has good judgments. That said, I think it's also the case that none of... I wouldn't bet that our... I wouldn't guess that our bet on New Incentives will work out, if I had to... If I was just trying to say, "Is it more likely than not that New Incentives will become a top charity?" I'm hopeful that they will, but that's really, really hard. And they're at the very beginning. And so I think what is most likely, is that we and they learn a lot about the difficulties in creating a new top charity. I certainly think they may, but I don't know. I'm pretty sure that that will be a function of the merits rather than something else. Yeah?

0:19:16 Speaker 5: So how many people are affected by the iodization of salt? How large is this population?

0:19:26 EH: Yeah, so I don't know... So the question is, how many people are affected by this salt

iodization problem?

0:19:30 S5: Yes, there's a lot of people, right?

0:19:33 EH: Yeah.

0:19:33 S5: Compared to other...

0:19:34 EH: So it's still a very large problem. There are... You can go on UNICEF's website, because UNICEF is the one that does the surveying for determining the extent of iodine deficiency, and there are countries where large proportions of the population are iodine deficient. So this is one of the main questions that we had as we entered this cause three months ago. I think the thing we thought was most likely to happen, is that we would find that this problem was more or less already solved, or at least solved in a way where additional money wouldn't make a difference. And as we've talked to people and actually looked at the data, we see multiple large countries where there are still significantly deficient populations. And importantly, the Gates Foundation has provided a lot of funding to, I think it's 15 or so, very high population countries. So India, Ethiopia, Nigeria are countries where the Gates Foundation is very involved, because clearly, if you can get a country to iodize all their salt, then you want to be working in the country with the largest population. But there are countries like, I believe, Madagascar, Angola that have been, do not have significant activity going on.

[pause]

0:20:57 EH: Yeah?

0:20:58 Speaker 6: You talked about AMF and how you'd removed them from your list because, from your top charities list because of lack of funding. Is there any worry that charities in that situation may be more aggressive than they should be in trying to spend some of their existing money in order to get... Because they know that if they do that, they can get back on the top charities list and potentially increase their money coming in.

0:21:20 EH: Yeah, so it's certainly... The possibility of creating an incentive for charities to spend money unwisely is a dangerous one. The simple metric of, "How do you spend money?" I don't think is a great one. We thought about this with AMF specifically. Based on what we know about them, I think they were very, very unlikely to take the message away like, "We should spend money more easily," as opposed to like, "Alright, we gotta find a way to get his done. We'll return money to donors" or something they were even talking about. I definitely think it's something we've talked about for other organizations that we know less well and it's certainly a big part of our update process. It's like, we not only want to see that the money was spent, but was the money spent well, did you do things... Going in, we have some understanding of what groups intend to do with additional funds, so, "Did this group spend money in a way that was consistent with what they said they would do, and does it seem to have been effective?" Those are the questions we're trying to answer.

0:22:27 Speaker 7: Relatedly, for AMF, one of the criticisms that occurred for being de-listed as a recommended charity is the idea that they need a large amount of cash in order to keep their pipeline fully funded, because they can't enter negotiations with a counterparty if they can't commit

to delivering on that negotiation, ultimately at the end. And there's a risk that all of the people that they're currently talking to could result in distributions, therefore they can't start talking to anybody new until those existing conversations are finished, either resulting in distribution or non-distribution. And so, because they don't have this buffer of cash, they don't have the ability to start more conversations and keep their pipeline fully populated.

0:23:19 EH: Yeah.

0:23:20 S7: So, do you have any comment on that?

0:23:22 EH: Yeah, so I think it's a really good question. Put very simply, is it appropriate for AMF not to be recommended because they need to have money in the bank to be able to negotiate with countries? And they're now... They've told us, this is in this page we've published within the last month, they're explicitly not trying to build out a deeper pipeline 'cause they're not sure that they would have money in the bank at the point they got there. I think the situation with AMF is like this: When they first got a recommendation from us in 2011, they completed their biggest distribution to date, which was about a million dollars' worth of nets. And then since then, they hadn't been able to finalize any distributions. And so, I think that there's some level of concern that AMF will just not be able to negotiate through the challenges of working with the governments and in the WHO and UNICEF, the way that is operating at a scale in negotiation that's very different than what they were doing before.

0:24:29 EH: That's why I think, this first distribution that they agreed to is definitely a good sign that there can be some movement there, but it's still a lot smaller than the amount of funds that they have. And so, I think what we really want to see is successfully negotiating a large distribution, using up, let's say, more than half of their funds. And then at that point we would say, we shouldn't ask AMF to always draw down its bank account every year. That would be a very unreasonable thing to expect and I think detrimental to their progress, but instead say, AMF now has something of a track record of being able to negotiate large scale distributions. Now, we expect that we'll have to wait and they'll be a timeline and a lag of money in and distributions out, and that is what we'd be ready to see them build up.

0:25:22 S7: So do you think there's any possibility of setting up some kind of trust, where AMF could draw that money if they had a distribution that they could commit to, but then they wouldn't get it unless they actually were able to accomplish it?

0:25:36 EH: Yeah, so Ian's asking this question of whether you could set up some vehicle where you accomplish a goal of providing AMF with money if it needed it without directly giving them the money. So I personally talked about this with Rob Mather, who's AMF's CEO, and so this is definitely something we're open to considering. I think this is a possibility if people were really interested in it. They could execute this on their own by giving to a donor-advised fund and telling Rob what they were doing or telling us what they were doing. I think it's also the case that a lot of... The dynamic we saw among our donor base... The donors use our research, last year in 2013, 'cause most people tend to give in the last month of the year. So something we saw is that a lot of people didn't give as much as they had previously, or didn't give at all and held back their charity for the year, or gave into a donor-advised fund, and I think a lot of this was explicitly with the intention of waiting a year to see what happened with AMF, and giving to AMF at the end of 2014, if they're still top rated and they need the money.

0:26:45 EH: But I do think it's an important point because often in charity, the organizations sometimes have a cash flow issue, where they literally need the money in the bank, and sometimes they just have a planning issue where they need to know that donors are there for them if they need it. And I think a lot of good can be done by donors just being very explicit with organizations about what they intend to do with their donations.

0:27:08 S7: It sounds to me like you're saying a lot of people are doing this anyway by waiting. So if AMF is unable to start a pipeline because they don't have the committed money and there's also all these people that are waiting to give, it just seems like all the pieces are there except the actual...

0:27:26 EH: Yeah, I don't know. The question with this waiting is exactly how much it adds up to and what's there. I also think just practically speaking, AMF knows what our position is. We're trying to be very transparent with them, and so this is something that we could be helpful to them by moving on, we would try to.

0:27:49 S?: Yep.

0:27:54 S?: Any other?

0:27:56 Speaker 8: So you mentioned a little bit about the challenge of trying to evaluate larger organizations, whether it's the Gates Foundation, NOAH, those kind organizations that are maybe doing very good work. But it's a [0:28:15] [REDACTED]. So I'd say, it's very difficult to evaluate the effectiveness. So are you... And then I read about sometimes, say, even larger organizations, that the money they have is still very small compared with the size of all the problems in the world. And then there are some who get experiments in terms chapter abbreviations amongst different kinds of organizations. So are you trying to see... Sometimes is it the case that you use knowledge, say, the Gates foundation or some other foundations who would take money from your donors, and then if they are very effective, but the scale of the problem is such that they would also need more money that you may somehow figure out a way to recommend them?

0:29:16 EH: I see. So I think your question is, is being large as an organization itself a sufficient reason to not recommend an organization, because certainly you could have a very large organization that is trying to tackle a problem that's even larger than its resources? We do face challenges investigating large organizations just because of the fact that they're very large. We want to understand the details of what's going on in an organization's activities and the larger they are, the harder it is to find a group that is able to communicate with us in a way that we can learn well.

0:30:00 EH: I think one of the specific ways that that challenge plays out, is the larger the organization, the more likely it is that they're more nervous about what a GiveWell report will say about them. It also means we often are talking to communications staff relative to program staff, which is just a sort of a specific challenge that you face when dealing with a larger entity.

0:30:23 EH: When I talk about larger organizations though, there's also this issue that we face which is that they implement many different types of programs. And so, if you think about, not to pick on UNICEF, but if you just take UNICEF as an example, they run health programs that I'm sure have great evidence behind them. They also run programs that I'm sure do not have as much evidence, like an education program or other types of programs. And it's really difficult to know

where the marginal dollar is going in that very large entity. And it's not just a function of the fact that they're large, it's a function of the fact that they're so diversified that you really are not... You're kind of throwing your money into the pot of this extremely diverse entity and you don't know where it will end up.

0:31:13 Speaker 9: So you talked a little bit about this morning donations to groups that [0:31:17] [REDACTED] countries to try to fill the evidence base, and I'm wondering are you talking about, or exploring restricted gifts or particular types of brand [0:31:29] [REDACTED] specific interventions or would this is, kind of restricted, just how the groups do that type of work?

0:31:35 EH: Yeah, that's a good question. So the question is, what kind of more specifics on what we're thinking about in these funding groups that are doing the randomized control trials. That's an area where we're really only focused on supporting very specific research. We went to them and said, "Broadly speaking, we would like to support additional research that adds to the evidence base about a program. Are there programs that you think are most worthy of additional research?" and they sent us some ideas. And then we also have some understanding of the evidence base currently, and so we said, "Here are some programs that we're particularly interested in." And so the programs now that were most interested in seeing more evidence on are programs where there's basically a single randomized control trial that found a great effect in one location and we're just really curious how generalizable that effect is.

0:32:28 EH: And so the programs are Incentives for Vaccines, which is basically giving women whose children are coming to get, women who bring their children to be vaccinated, a bag of rice or something, as an incentive to get their child vaccinated. This had great results in a single trial in India, hasn't been replicated as far as I know. Incentives for Migration, which is giving people some money and saying, "Hey, you can seasonally migrate within this country and earn a lot more if you migrate from rural to urban areas". And then there's an eyeglasses study in China where giving out eyeglasses to students resulted in significantly higher test scores, basically, so that's another one we're considering. And then finally, there are some... There's been a study of commitment savings devices in the Philippines, and so that's sort of the fourth one that we're talking about. I think a lot of what we'll end up deciding what we move forward with is largely a function of what the organizations think they can pull off, like can they find a partner who wants to implement it? Is there an academic who's interested in researching this program as part of their academic career? So that's what we're looking at.

0:33:31 EH: Cool, so I want to pause here and just move on to the second half, the GiveWell Labs side of things, and so again, I'll just kind of give a brief overview of what we're doing and then we can pause for questions and also leave some time for any other questions that people have at the end.

0:33:48 EH: So just one quick note on GiveWell Labs is I spend more of my time personally, on the things we've already talked about, so I'm not as involved in all of the details of GiveWell Labs, and so it's possible that I won't be able to answer some of your questions, but I will certainly try.

0:34:04 EH: So, GiveWell Labs is the more open-ended side of GiveWell. It's where we're dropping the criteria of, dropping the evidence-based criteria that is so important to our traditional work. And we started GiveWell Labs a couple years ago for a couple of reasons. One was that the types of recommendations GiveWell's traditional work was making, didn't seem well-suited to the

very largest philanthropists. These are people who could start organizations from scratch if they wanted to because they have sufficient funds. And so, we wanted to have a research product that was helpful to that type of donor. And working closely with this foundation, Good Ventures, which is the sort of primary short-term consumer of this research. The other reason is that GiveWell's mission is to find the best giving opportunities we can, and I think those giving opportunities could easily be high-risk, high-reward opportunities, as opposed to strong evidence of impact opportunities and so, that's why we're looking at it.

0:35:11 EH: One of the ways that GiveWell Labs differs most concretely from our traditional work is the traditional work focuses on the unit of the charity. Is this charity good implementing a specific program? And GiveWell Labs focuses at the cause level, where we're instead asking, "Is this a cause that we should get involved in?" There's a few reasons for this. One of the main ones is that, again, if you imagine being able to start an organization, you don't want to be limited to the organization that exists. You want to know about areas that may need to have new organizations created. And I think this is actually something that philanthropists, I think, have done reasonably well over the years, finding an area that is getting insufficient attention from the world and bringing more attention to it.

0:35:57 EH: The other part of this is without the criterion of evidence base, is this keeping you go on? Having more context in an area where you really have committed to a cause, understand a cause, I think puts us in a better position to separate the better from the worst proposals that we see.

0:36:15 EH: So generally, on a substantive level, we're sort of separating the GiveWell Labs of world into four categories of causes, two of which we haven't put much time into so far, though we've put some, and I can talk about it if anyone's real interested, scientific research, and foreign aid. Obviously, foreign aid is in some ways covered by GiveWell's traditional work, but we're also looking at it for Labs. So both of those are our areas where we haven't done very much yet.

0:36:42 EH: The two areas where we have put more time into are politics, which is influencing laws, regulations in the US. We've primarily focused here so far, though that wouldn't be something that was necessary long-term. And then also global catastrophic risks; these are things that have a very low probability of some very bad outcome for humanity. And so, where we are right now with both of those causes, and I'll go into more specifics, is that in 2014, our goal is to commit to a few causes in those super causes or those categories by the end of the year, in order to put ourselves in a position to... We think committing is necessary in order to learn about what we can do in the cause and what the opportunities look like.

0:37:38 EH: So we've done a lot of work that is relatively high-level, comparing between causes, and this appears on our website, but a challenge that you have in philanthropy is it's not like all of the good ideas are just out there and you can kind of go through them all and compare them to each other. Instead, you often need to commit money and staff time to an area and say, "We are a funder here," in order for the best ideas to come of the closet and come to you and be assessed.

0:38:08 EH: So what this commitment really means for us is hopefully by the end of the year, we're in a position where we can say, "Here are some areas where we're committing a significant amount of full-time staff, and something like \$5,000,000 to \$10,000,000 a year for a few years to this cause." and we hope that that time of commitment enables us to get these learning benefits from these areas. And I say learning benefits 'cause these are not the causes that we have decided are

necessarily the best, long-term. We're trying to pick causes that certainly now seem very promising to us, but are ones where we expect that the learning that we'll have over the next three years will put us in the best position to learn more in the future.

0:38:49 EH: So now I just want to go through some of the specific causes that we're thinking most about. So the first area is politics, and there we've gone through a very long process of talking to what I would call people with general policy knowledge about which areas seem most promising. Also then going through a whole list of causes to assess them on what are our key questions I'd say in GiveWell Labs. And so, these questions are, "How important is the cause?" meaning if you were to have an impact, how good would that impact be? What would it mean in humanitarian terms? How tractable is the cause? Is this something where you could realistically imagine having this impact? And then finally, how crowded is the cause? Is there already significant funding in that area, with the notion that to the extent that there is more funding that likely reduces the marginal value of additional funding. Now, obviously all of these sort of interplay with each other. A cause could be crowded in a certain way that's low impact, very uncredited at a higher impact, but conceptually, these are the types of questions we're trying to ask.

0:40:05 EH: So the three areas... I'll go through the three areas of politics that we've spent the most time on so far and are sort of closest to committing to. So one is criminal justice reform; this is an issue that was highlighted to us by... This is in the United States. And the goal here is reducing prison population and increasing safety. And this was highlighted to us as a cause that is particularly tractable at the moment. It sort of stands out among other causes for the potential for making a difference and the reason is that for a long time, I guess the left, has been very interested in reducing the prison population for humanitarian reasons. The right has been known, I guess, more for tough on crime, increasing surveillance and increasing prison population. But more recently as states have faced budget crises, this has also become an issue of where states can save money by reducing prison population. And so it has something of a trans-partisan consensus perhaps.

0:41:09 EH: What we have largely done so far, is fund a couple... And when I say "we," again, this is Good Ventures providing funding, fund academics who are working on this area. One of them is a guy named Mark Kleiman. He has some of the more, we would say, insightful or informative views on criminal justice reform that we've seen. One of the main ideas that he's well known for is the idea of swift in certain sanctions, which means trying to deter crime not by increasing the severity of the punishment that is meted out, but instead making sure that people are swiftly and consistently punished for that crime. And there was a randomized controlled trial of this program. It had pretty promising results. The idea is we don't need to have people in jail for as long to get the same public safety goals that we have.

0:42:00 EH: Another thing that he is well known for is promoting the idea that we need to better regulate alcohol. There's a lot of talk about legalization of drugs, and he has worked with Washington state on rolling out their legal pot initiative. But on the other hand, he's saying there's other substances that are abused, alcohol is one that causes a great deal of crime, and we should do what we can to better regulate that area. The other academic that we've supported is a woman named Angela Hawken at Pepperdine University, and she's trying to provide states with the tools to evaluate their own programs rigorously through randomized control trials. So we've supported this initiative called BetaGov, that's Angela Hawken's group.

0:42:45 EH: And then finally, we're kind of talking a lot to a relatively large public charity that has a criminal justice program. This is more explicitly political; they work pretty closely with state governments. They take a more technocratic approach to the work that they're doing where they're finding small ways to tweak laws that they hope can reduce prison populations while maintaining public safety, and so we've also talked a lot with them and may be providing funding to their efforts.

0:43:16 EH: The other two causes are definitely on the side of potential for very large impact and probably not particularly tractable. So one is a cause we call labor mobility, which is trying to enable people from poor countries to migrate to richer countries. This is something that Michael Clemens, who's an economist at the Center for Global Development, he'll point to this as the best poverty alleviation program that we know of, and he has an estimate that open borders would lead to a 2X, two times, increase in world GDP. We looked at the numbers, I don't think we would stand behind that estimate, but I certainly think the notion behind it is sound that, you take someone who's living in Haiti, and they move to the US, and if they can find a minimum wage job that has an unbelievable impact on their standard of living. This is definitely a controversial issue.

0:44:13 EH: It's one where there's a lot of battles going on right now in government on immigration reform, and so the interest we have in it is that we don't see others making the humanitarian case for more open immigration laws, and we don't see any philanthropy to speak of on that side of the debate, and so that seems like a potential interesting opportunity for impact.

0:44:40 EH: The final politics area where we're kind of most seriously considering going into, and there's a longer list of causes on our website, is macroeconomic policy. So this is an area where global recessions cause a significant amount of financial harm in the rich world and in the poor world. And if recessions could be better managed, you could a very big impact on humanitarian outcomes. It's definitely... So there's a lot of things that are unclear in this area. It's unclear how crowded it is, first, I mean there's academic research done on macroeconomics; the groups like the Federal Reserve fund their own research. But there certainly is no major philanthropist that has macroeconomics as their policy area. And so that, in some ways, is a sign of a lack of un-crowdedness.

0:45:33 EH: It's also the case that... So then there's also this question about, what would you do in this area? I think this is something we're still trying to learn more about, but certainly in the conversations we've had with economists, a couple things have come out. I think you could potentially fund more practical research that helps us better know what to do in recessions. It's also possible that you could fund improved communications, such that the government were able to undertake the policies that would be likely most effective. Obviously this is a very controversial... It's a challenging environment. There's no notion here, or with labor mobility, that these would be particularly easy. They're particularly hard for sure, but the lack of philanthropic involvement and the potential for very large impact, makes them appealing to us as causes to investigate further.

0:46:26 EH: The last area I just want to touch on briefly is this area of global catastrophic risks. So these are things... I think the example that people might know most about is like an asteroid could hit the earth and that could cause the extinction of humanity. It's very unlikely that this happens, but if it did happen, that would be really bad. This is the type of activity that philanthropy seems particularly well-suited to supporting. This is not something that a market actor will take care of. It's possible that governments would take care of this. I think in the case of asteroids, it's significantly

less promising than it might seem on its face, because NASA has done a pretty good job of this over the last 20 years, of identifying the larger asteroids that are out there.

0:47:08 EH: But there are causes that seem to us particularly scary and don't have much philanthropic involvement. And so, these are biosafety, so then this would be like bio-terror or flu pandemics, and trying to find ways to reduce these risks. There's a lot of government funding in this. There's about \$6 billion a year on biosafety, but it tends to focus mostly on bio-terror, and I think our understanding is less on things like pandemics. And so conceivably you could support activities like improved research on the potential magnitudes of the impacts, or better ways to mitigate the potential for a pandemic.

0:47:50 EH: Another area is geo-engineering. So this is coming up with technology that allows us to cool the climate in the event that we have a sort of disaster or climate change scenario. This is something where, again, very little philanthropy is there, less than \$10 million a year. It's something that is particularly controversial because it... By funding this type of technology, you might reduce the incentive for carbon emission reduction. It's also possible that by developing this technology, a country could get hold of it, implement it, and then cause some sort of global war over the technology and so, a lot of... These are the questions that we're thinking about as we look into this cause, and one angle on this cause that we find particularly interesting is an angle of management, or safe management, of emerging technologies, which also fits well with the biosafety area.

0:48:50 EH: And then the final one is artificial intelligence risk, and this is something that I think most people, certainly I did this, sort of laugh it off the first time you hear of it, like the robots taking over idea. But from the conversations we've had with people who know more than us, it certainly seems like something that could be the biggest technological change of the next 100 years, let's say, and it doesn't seem that any philanthropy, or significant philanthropy, is really devoted to thinking about how we should best address and manage this risk. And so that's a final area that we're looking into. I mean, we've kind of now looked at what we think are most of the global catastrophic risks, and these are the three that we're most worried about, and most focused on, or at least, tentatively most focused on for the time being. So I want to pause there and see if anyone has any questions on that on GiveWell Labs.

0:49:42 Speaker 10: Could you go into a little bit of why artificial intelligence is up there so high on your list? Like, some of the research that you've done that makes you put that as one of your top three things to focus on?

0:49:54 EH: Yes, so the question is, why artificial intelligence? I don't think I'm going to be able to give a great answer to that 'cause I don't have the right background. I would say that we've looked at a list of the risks that kind of came out as global catastrophic risks. And so the ones I haven't... I mentioned asteroids. The ones I haven't mentioned so far are, these are always fun to talk about, giant super volcanoes. So for those of you who don't know, Yellow Stone national park is a giant volcano that goes off every six hundred thousand years. If it goes off again, we're in big trouble. That is one where it's not clear what if anything could realistically be done. I guess you could fund research into it. I think also just intuitively, we're more discounting the natural catastrophes relative to the human created catastrophes with the notion that, "Hey, we've made it so far, so it's less likely to be catastrophic as a man-made disaster." And then another one is nuclear safety. That's one where our impression... We have a page on basically all of these on our website now, but it's one where there is significantly more philanthropic involvement.

0:51:05 EH: I don't... I think AI is the one that we know the least about, but it's certainly one that of the full list of areas we looked into is one that people pointed to as scary. And you'd have to talk to, I mean we could talk to someone else at GiveWell who could give you a more concrete answer about it. Yep?

0:51:26 Speaker 11: What about the list, another list, genocide? I mean, there's other government agencies in this country that recently, there was a coalition set up for the event of genocide. But I mean, could you consider that?

0:51:46 EH: Right, so the question is, does genocide fit in? Where does it...

0:51:51 S1: Some countries don't even say, they don't acknowledge the word genocide.

0:51:56 EH: Right.

0:51:56 S1: So it's hard to prevent it or set up enough conditions to prevent it, just 'cause the word is so political.

0:52:03 EH: Yeah, I mean I think the way that I would... The way I would describe the global catastrophic risks are things that could themselves significantly set back humanity, like, broadly defined. So that's why something like an asteroid that wipes out humanity, or some significant proportion of humanity, or some pandemic that does the same, that's sort of what fits into this category, where they're just very, very unlikely outcomes, but they would be so terrible. I think genocide, in that issue, we would put more into a different category, like a category of conflict between states and what can be done there, and that's... I talked about these four super, these four categories of causes. Another category that I didn't mention 'cause it's kind of a one-tier lower down on our current priority list, is things that involve state-level interaction or state-level negotiation, and this is something where it is... The next step that we want to take there is get a better understanding of what philanthropy has or has not accomplished in the past in that whole enterprise because we're not... We want to better understand what philanthropy has done to get a sense of what it could do when you're talking about state-level discussion.

0:53:40 Speaker 12: There are already a large number of think tanks that conduct economic policy research and policy advocacy, and I was just wondering what do you... It doesn't... It seems to me like they are already saturated. What do you think that GiveWell would have to have?

0:54:03 EH: And so you're asking about like in macroeconomic policies specifically?

0:54:06 S1: Right.

0:54:08 EH: Yeah, so the question is, isn't this area already saturated? So again, unfortunately, I don't think I can walk through the specifics of this. My understanding is that today, the research being done, is not sufficient, given that it's not the type of research that could be done. And something that you could look at is the... On our website we'll post conversations that we have with experts. We've talked to a few macroeconomists who will say like, "Oh, this, this and this isn't being done." I think there is something happening now. I think one of the major think tanks is setting up its own division, explicitly focused on macro economic policy, and it's certainly possible

that that entity could resolve the issue that we have. It could solve the problems so to speak, but that's something that we're watching. I don't know whether that will happen or what they'll end up doing.

0:55:09 EH: The final thing is I do think that GiveWell's role would not be to do all of this research ourselves, but fund the people who are doing the research or identify people to do, nudge them in the direction of doing a different type of research. And so I think the model that makes more sense is one where you say... The current research that is available or the current activity that is out there is not ideal and then we're trying to... We could provide support to create an entity that is doing that type of work.

0:55:42 EH: Finally, this is sort of my own general impression that a lot of the, a lot of the think tank work often is like particularly partisan, and I think that leads to a certain type of analysis, and I think a goal would be to create a less, a type of analysis that's less focused on one particular agenda. I guess everyone probably says that exactly.

[laughter]

0:56:10 Speaker 13: Can you talk a little bit about where you are in learning the lessons of how to evaluate the whole space of causes and how to figure out which causes are actually the most important?

0:56:19 EH: Sure, so the question is, where we at in learning about how to learn about causes.

0:56:25 S1: Yeah.

0:56:26 EH: Yeah, I think what we're trying to do... The thing that we think is really important is that we can't just stay at the highest level for too long just evaluating things because we've learn a lot about what it means to evaluate a cause by getting involved in that cause. So in the case of criminal justice reform, it's a really good example where getting into the details of the cause and meeting with, learning from people like Mark Kleiman and Angela Hawken, and then this other public charity, these are all things that we wouldn't have been able to do from afar. And so, I think a lot of what we're learning is that there's a... It's necessary to go both broad and deep at the same time, even when... We're at a point now where we're ready to go deep on certain causes even though we're not at all confident. They're like, "These are the best causes," or that, "We won't find better ones in the future."

0:57:23 EH: This is a lot of the way that GiveWell has operated historically, where we've... From the very start, we wanted to find charities we could recommend, even though we recommended these organizations after only three months. We knew these were not going to be the very best charities we could ever find, but by getting ourselves to sort of the end game of our research, we were better able to figure out what we needed to know at the higher level.

0:57:47 EH: Another part of this is something I sort of alluded to this idea of being ready to commit to a cause such that you're able to learn about what exists in that cause. This is something that... Another thing we've tried in our GiveWell Labs activities, and this thing we'll be writing about a little bit more in the future, is we hired a consultant to go out and survey all the big malaria organizations, about what projects they can't fund but would be incredibly impactful. And we got

back all these proposals from these organizations. We ultimately decided that we weren't ready to move forward on Malaria at the time, but one of the things we found is that when you go to an organization and say, "Can we solicit proposals from you?" they generally only want to talk to you if they have an expectation that they're likely to get funding.

0:58:41 EH: And so, this is very challenging. We tried to be very explicit about the likelihood of funding, but you can't just sit back and, I think that nonetheless that type of project that we did in Malaria would not be feasible in 20 other global health causes because it just wouldn't be the way that the sector operates. So part of what we're learning is you want to go to this deep dive to learn about what questions you have, but you also need to do the deep dive in order to get the information necessary to really understand what a cause looks like from the inside.

0:59:15 Speaker 14: So related question, the other side of the question I guess. As you guys have grown, or as you interacted both with charities and you've seen how other charities have interacted with other donors, how much do you think that you're changing the model? 'Cause I think that's one of the things that GiveWell donors, or potentially excited about is the fact that you guys are very focused on evidence-based analysis and you're doing deeper dives. Do you see that model being adopted by other donors, and do you see charities responding to that?

0:59:47 EH: Yeah, I think we see... We see some movement on the charity end of an interest in being more transparent. Certainly, that's grown a lot, as the GiveWell's money moved has grown. Money moved is the money that goes to our recommendations as a result of our research. I think there's definitely more of that. I don't think I would point to it as... Changing charities' behavior is not something, I would guess, is going to be the way that GiveWell has its biggest impacts. I think it's actually on the front of changing the model. I think that we're... The place that I think we're more likely to influence is, "Oh, no, we can't do this directly," but is funders.

1:00:33 EH: Basically, there's very little reason that funders shouldn't be more open with the work that they do and in some ways it's very... It's quite unfortunate that we're trying to look at all these causes that others have looked at in the past and we're essentially starting from scratch in looking at them. And so I think that's something where I hope we can have influence. I hope that people start to do things differently. Certainly, Good Ventures is following us in being very transparent. Hopefully, they can also serve as a model. And I think you even see some movement on this with the Hewlett Foundation has recently started publishing a lot more of the things that they're thinking about. And if I had to guess, I'd say that is really a very realistic way that you could create a much richer set of information about philanthropy than exists today. Yep?

1:01:24 Speaker 15: So I'm going to... To what degree, if at all, are you guys engaging in direct advocacy with philanthropy organizations and other groups of donor networks were already engaged in some degree of advocacy for transparency among donors?

1:01:41 EH: Yes, so the question is, to what extent are we actively pushing this idea of more transparency among funders? We're not proactively doing this in any major way. I think what we see... The best thing I think we can do right now, beyond obviously talking to people when they ask, and we've had conversations with groups when it has, they thought it would be helpful, is just being a model of how this can be done, because we've come up with a lot of processes around our traditional work and how we're transparent about it and how we protect confidentiality in the charity's interests and our own interests for transparency. It's more challenging, certainly on the

political side, because you start to have a situation where you literally have someone who's on the other side, who you're against. And it's not clear the extent to which you can share it.

1:02:34 EH: So I think these are things where there's still room for us to learn about what we can do, and then just by ourselves trying to be a transparent funder, showing others what they can do. The final thing we're doing that is just kind of related to this idea is this... So then we call it "The History of Philanthropy Project" where we have hired some consultants to do a more historical journalistic assessment of past philanthropic successes, and this is very explicitly trying to go back to foundation activities or successes of 20 years from now, 20 years ago, and saying, "How well did this work? What did they do?" And the fact that it happened so long ago makes people more amenable to talking openly about their activities.

1:03:32 EH: Alright, last question. Or two questions, and then we'll wrap up.

1:03:35 Speaker 16: Mine is right off the bat. This one is just a clip, I just want a clarification. When you were talking about being against someone in this process, who were you talking about being against? What organization are you necessarily pitted against? Is it another organization like yours that's trying to, in a way, do the same things and they're trying to get some proprietary information?

1:03:57 EH: Yeah, so let me give you an example of what it could be. It'll be a cause I haven't talked about much so far, but I think something that I would say we're looking for are causes where there'll be pretty broad agreement of the humanitarian benefit of an issue, or at least we've written... People ask often, "Oh well, GiveWell's research must be a function of GiveWell's staff values." and this is true, and so we try to write publicly about what those values are. And so, our values would say this cause has humanitarian benefits. That's an area where we're looking for. And then you might have some opposition that is not interested in the humanitarian outcome. They're interested in their own personal benefit.

1:04:41 EH: So a good example of this is industrialized agriculture or factory farming. Basically, you look at the situations of these animals and I think the conditions that they have to go through are incredibly bad. The people on the other side of trying to create laws that prevent that type of activity tend to be the industrialized agriculture lobby, and so that's the type of group you're fighting against. And they're not a group that is fighting against it for humanitarian reasons or fighting against it because that's the activity they undertake.

1:05:15 S1: Okay.

1:05:18 Speaker 17: So you had alluded to this, I think in a blog post, the idea that GiveWell might de-merge into GiveWell's traditional organization and the GiveWell Laboratory organization.

1:05:29 EH: Yeah.

1:05:30 S1: Do you want to say anything more about that?

1:05:31 EH: Yeah, sure. So the question's about the sort of long-term vision of GiveWell Labs and GiveWell, and how they'll be intertwined in one or two organizations. So these are clearly... In many ways, these are very different research products. They're very different activities, and so in

many ways it seems like they're being separate makes a lot of sense. We're also cognizant of the fact that getting involved in something like politics can cause harm to the sort of GiveWell brand, which is associated with more or less like unbiased objective assessment of evidence.

1:06:11 EH: And so in many ways, the long-term vision of these just being two separate organizations, both of which do research, both of which are transparent, that is very attractive in certain ways. So I'd say our default expectation is that is what will end up happening five years from now, let's say. That said, if I look back five years, it would be very hard to predict where we are today, and so it's not unlikely that something changes. The main reason we haven't done this thus far, is that right now our staff is more or less flexibly allocated between both sides of GiveWell, and it would just create a lot of problems for us to try and do the research we need to do with an explicit staff split, and so for now, we've maintained ourselves as one single organization.

1:07:03 EH: Alright, cool, so I think, about time to wrap up. I just want to say I really appreciate everyone's coming out to this very non-standard charity event. We, all of GiveWell staff, really appreciate the critical engagement that we get with our research. It's a lot of a reason that we want to put all this out there is to get this type of feedback and questioning and we don't take that lightly. So thanks so much for spending your evening with me.

1:07:30 S?: Thank you.

[applause]