

**GiveWell San Francisco Research Event, November 18, 2015 –  
Open Philanthropy Project**

*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.*

00:00 Holden Karnofsky: A little update on the Open Philanthropy Project. I know we went a little long on GiveWell because there's a ton to share and a lot to sink your teeth into and Open Philanthropy. I'm going to go on the shorter side. I'm going to give a very brief overview of where we've been so far, and I'm not going to give all the background. One thing to say about the Open Philanthropy Project is it's a new idea, it's a new thing, it's a new project. It's actually been going for a couple of years, but in terms of the amount that we still have to go and the amount that we've developed, it's much more like GiveWell was in 2008 or 2009, and it doesn't have its own website. And so, it is very hard to learn about right now. And one of the big things we're working on is trying to get a website up for the Open Philanthropy Project. That will make it a lot easier to get to basics, a lot easier to follow it. But since I see a lot of familiar faces and since I've been over the basics of what it is before at these meetings, I'm going to try and be fairly efficient and give the basics and then talk about what's new and then take questions. And if you run out of questions about Open Phil and you want to go back to talk about GiveWell, I'm happy to do that too, although I'm not sure it'll go as well without Elie here, but give it a shot.

01:06 HK: The Open Philanthropy project... Let's see how quickly I can do this. I haven't really tried this before. The Open Philanthropy Project and GiveWell have essentially the same mission statement, which is, you find outstanding giving opportunities and you publish the full details of our analysis to help donors decide where to give. So, it's this idea of bringing open public analysis to philanthropy and looking for the ways to accomplish the most good with money. But they have different audiences, and that causes them to be different in almost every other way. So, GiveWell is really targeted at individual donors. It's targeted at kind of the person that Elie and I were when we started off which is, "Hey, I work at..." Let's say a hedge fund. I have, let's say, ten hours or five hours or one hour or five minutes, to decide how much to give this year, and I think I want something that's verifiable, something that's been audited, something that's been vetted, something I can believe in, and that kind of leads to the direction of proven cost-effective scalable organizations that can take in a lot of money and do a lot of quantifiable good with it.

02:06 HK: Open Philanthropy Project is more directed less at the me and Elie of 2006, and more of the Kerry and Dustin of now, which is, "Hey, I'm a major philanthropist. I'm going to give away let's say several billion dollars. Instead of having an hour to think about it, I have like 50 years." So, it's a different problem, and you have different options, and I think therefore, there's different things you want to do. So as a donor like that, you can build up your own expertise in a way that individuals can't. You can build up your own staff which I think is an extension of building up your own expertise in a way that individuals can't. You can get much deeper into issues, you can help create organizations, you can help transform organizations. You're not stuck with what organizations already exist. You're not stuck with what's verifiable because you can either become or hire, and build trust in, someone who really knows an area and knows all the soft things that would be hard to uncover by looking for randomized control trials.

03:03 HK: So, I think it's a different set of things and I think you end up being much more interested in very long-term stuff, very high risk stuff, maybe doing funding things that may not look good to anyone who's not totally in the know, and may take 30 years before they generate anything of note, but if they work out, it would be really incredible. If I were to refer to historical examples of this, the Rockefeller Foundation was once funding research on how to improve seeds and get them to get higher yields, and this was not the hot topic of the day, and it was the opposite of proven. It was like trying to create a new kind of seed. But this research has been credited with playing the major role in launching the Green Revolution which pulled multiple countries basically out of poverty, turned India from a wheat importer to a wheat exporter, ended a famine and basically, is considered one of the biggest humanitarian events of ever.

04:01 HK: And so, that's the kind of thing you'll be going for. Another example would be the birth control pill, the combined oral contraceptive. That was something that was not at the time, the research to get that thing together, that money was not coming from the government. That came mostly from a woman named Katharine McCormick who really believed in this, and wanted to try to fund it, and it was not the hot thing to be funding at the time and not something that you could prove was going to work with randomized control trials. So, that's more the idea, and we're looking for stuff like that while realizing that unlike a lot of foundations who I think take it for granted that that's what you want to do as a major donor, I think we remain a bit agnostic. We say, "Hey, if we can find stuff that we like better than GiveWell's top charities, then we're going to try and do it and we feel like we got to look for it because as a big donor, you're in a position to do things like that. But if we don't, there's no reason you can't just decide to give to top charities instead."

05:00 HK: And just to talk about the mission of Open Phil a little bit, it's not a consulting group for Kerry and Dustin. It is a partnership. We work on it together. We have a lot in common with our values and we've been working together for a very long time, so we feel that we're able to work together as a team on this, and the goal is not just to figure out what to recommend a good venturist should do, but to create a public set of very thoughtful recommendations that can inform a lot of donors. And certainly, individuals who decide they want to give in to that pool will be welcome to do so, and will be able to do so in the long run. But I think a lot of the target is also just future philanthropists in Silicon Valley and trying to build up the research to the point where we can say, "We've thought about the question about how to accomplish good with your money harder than anyone, and we're sharing more of what we do with it than anyone. And so we think we should have a major influence on how people give." That's what the Open Philanthropy Project is. In terms of what we do and in terms of the contrast of GiveWell, GiveWell looks for charities, organizations that already exist and the criteria are proven, cost effective, and scalable.

06:06 HK: Open Philanthropy Project looks for causes. So, because we don't want to be stuck with the charities that already exist, so, we look for problems in the world, issues, things like criminal justice reform, or climate change, or bio security. And we try to choose the problems to work on, or the focus areas that we think combine being important, neglected, and tractable. So, important means they matter a lot, they affect a lot of people. Neglected means we look at who else is working on them and we don't want to be the 20th major philanthropist in a field if there is another equally important one where we could be the first. Then tractable just means that we can see a path forward. So, we don't just pick things that would be really great and then nobody's trying. We want to feel like they could be done too.

06:50 HK: So, that's how we operate and in terms of what we've done so far, it's kind of a bigger... It's more of a boil the ocean project, probably than GiveWell. GiveWell figured out its criteria pretty early on, and there's a small number of charities that fit them. I think this is going to be a long project of gaining enough basic knowledge about an area to start thinking about who we want to hire, and then hiring those people, and then building relationships with those people to the point where we really have the expertise on staff to do things that I think we couldn't do as just individuals. So what we've done so far is, we have put most of our work so far into cause selections, so that's choosing what problems to work on. And we had a big process that was most of last year where we took every US policy issue we could work on, and why we started there I can talk about, but one answer is that we had to start somewhere and we haven't looked at everything yet.

07:44 HK: But we looked at every US policy issue we could work on, and we looked for the ones that affect a lot of people, that don't already have a ton of attention, and where we could imagine some political wins. And we kind of declared a bunch of focus areas based on that process, and did the same for global catastrophic risks. So, we looked for the things that could really disrupt global civilization, and we looked for things that are important, in the sense that things like climate change or pandemics, things that are important in the sense that they could really threaten civilization and it's realistic for them to do so. Neglected and also tractable in the sense there's something we could do today to reduce the risk. So, we spent last year basically choosing causes, and what we've been doing this year, well, a lot of what we've been doing, and I can get a little bit later to everything we've been doing, a lot of what we've been doing is then trying to hire around them. And so, one of our priority causes was criminal justice reform. That was the first area where we tried to make a hire.

08:42 HK: We spent several months working on our job description, circulating it, interviewing people, really thinking through what we were looking for and why. And we ended up hiring Chloe Coburn, formerly of the ACLU, and she is leading our work on criminal justice reform now. She started in late August, and basically what we've been doing with her is just getting to know her and working with her on some early grants. So, where we're soon going to recommend the first couple of major grants that are under Chloe's criminal justice regime. It's going to be two grants for about \$3 million total, one of them to an organization that's not public yet, that is, I think, going to be a major national player in criminal justice reform, and another one to a more specific, state specific campaign that we're excited about. And in that process, we've just been very intense about these early grants, about understanding how each other are thinking and understanding the case for and against these grants and building the trust, so that we can eventually head in a direction where we feel like we're synced up enough with Chloe that she can make grants we feel great about without putting a ton of work into them, and without putting a ton of thought into them, and only knowing that they come from Chloe's recommendation.

09:58 HK: So the basic model of the Open Philanthropy Project is pick great causes, then pick great people to work on them, then build trust with those people. Those three steps are really, really hard, and if you can get them right, the rest of it is really, really easy, I hope, which is that those great people working on those great causes don't need a ton of oversight forever. They can do great things with a lot of freedom if you put a lot of trust in them. Other hires we've

done, so, farm animal welfare is a cause that we declared a focus area. And we conducted a search and Lewis Ballard joined in late October, so very recently, he's working on his first couple of grants and has basically made the case to us that our priority for farm animal welfare should be corporate campaigns.

10:41 HK: So basically, he believes that trying to get people to go vegan or eat less meat as a way to reduce animal suffering is not nearly as promising as trying to go after very high leverage points, corporations, things like Tyson Foods, or McDonalds, or Costco, and think about what standards they're imposing on the chicken and on the other meat they sell. Because if you can get Costco to declare that they're not going to have chicken from battery cages anymore, that's a lot of chickens that you got at one fell swoop improved their lives. And how that cause and why that cause I think would be another conversation that I'm skipping over for now, 'cause we have talked about it before. And we are going to have... We do have information up about it, but that organization or the information is going to improve.

11:27 HK: Another focus area we have is biosecurity. This is in the global catastrophic risks category. And this is trying to prevent a pandemic, either a natural pandemic that you could imagine could be as bad as the Spanish flu that happened at the end of World War I that killed more people than World War I, or in some future world where synthetic biology improves, you can imagine that pandemics could get a lot worse than they've ever been before. And so, when you're thinking about things that could really threaten civilization as a whole, that's probably at the top of the list. There's a fair amount of government attention and very little philanthropy.

12:01 HK: And as we've gotten to know the space, we've gotten a feel that there's actually probably not enough attention on planning for very unlikely but incredibly bad worst cases, such that the kind of... The product of the probability and the badness would actually be something pretty substantial. That scenario we've basically been doing a search since... Kind of all year but really been intense about it since June or July. And we're not there yet. We have one person we're doing a trial with. But we are really, really thoughtful and careful about who we hire for these senior roles that are going to have a lot of autonomy and we don't want to make a hire if we don't feel great about the hire. And actually, even if we love a cause, if we can't find the right person, we're just not going to do it.

12:40 HK: So that's why we have a search ongoing and we have a trial hire ongoing. And then, another focus area is potential risks from advanced artificial intelligence. So this is another focus area that could be a whole conversation, why we're interested in this cause. But certainly I believe that if you're looking at things that could change the world a lot in the future, advances in artificial intelligence is a big one that has been really live in the last few years, could be a really huge deal in the next couple decades. And there's a lot of excitement about making AI faster and smarter. When you look at the people who are trying to plan ahead for things that could go wrong and how bad it would be if you had something that was smarter than, basically, smarter than humans in a lot of different things, how bad that can be and how you might prepare for it, I think that's a surprisingly thin field where there's very little going on and where we believe there should be a lot more thought.

13:31 HK: Again, it's a very early stage and in some ways you should think of it as probably useless to work on this today. But if it is useful, I think the value makes up for that and I think there's a decent chance that it is useful. So, on that front, we are working on a trial basis with Daniel Dewey who is basically helping us put together a landscape of the field, understand the groups that are already out there, Miri and others, and figure out where the biggest gaps are. What we would do if we entered this space, what the most important areas are that nobody's already funding, and then we'll see where that goes. And so that's... Really, a lot of what we've done this year has just been put a ton of work into trying to make hires in those categories and then other stuff we've been doing. Alexander Berger has been helping with the hiring and recruiting on policy, but also he works on some causes that are not very well developed fields and where we don't feel we'd be able to make a hire and where there's not a ton of funding opportunities.

14:30 HK: So he's also been... He works on three different US Policy priority causes. And then, we've also got stuff that's in earlier stage. So we have focus areas when it comes to US Policy. We have focus areas when it comes to global catastrophic risks. That means we've looked at all the options and we've kind of asked the questions and we've chosen what we think are the ones most worth working on. But if you want to talk about scientific research, all the different ways you can approach scientific research as a funder, and all the different pieces of advice we've gotten about what the best thing to do is, I think we're still pretty far away from having the right kind of staff to sort through all that, and from sorting through all that and making our priorities. So it's something I've spent a good deal of time on this year, but I think that's a huge project and I think we're not where we want to be. But we are working on a trial basis with one senior advisor, Chris Summerville, who I think having someone in that role would be a big help to us.

15:25 HK: I've skipped over a lot of stuff, but I'm trying to give a basic picture of what Open Phil is and where we've

gotten so far. So with that, I'll take questions, and feel free to ask about anything, even basic stuff that I skipped over.

15:39 Speaker 2: Have you done any evaluation of... Basically what you guys have done, which is philanthropic research [15:48] \_\_\_\_\_ more research into cost-effectiveness of other sectors?

15:56 HK: Sure, have we thought about the cause of cause prioritization itself? In some ways, yes, we thought about it, we're doing it. In other words, I would say we're not prioritizing that right now. I would say that we actually want to go non-meta before we go meta. I think what we're trying to do is build a framework for choosing causes, 'cause we don't think there is a good framework out there right now. We don't think there is an organization doing this in a helpful way right now. And so, we want to be that and we want to understand how it works. And then, I think we'll be a lot more informed and a lot better positioned to talk about what kinds of things more like that need to exist. So, for us, the priority has really been build a process and build a way of attacking some of these direct issues that I think you can do a lot of good, and the meta stuff is not at the top of the cue right now.

16:46 HK: I also think some of the meta stuff is... It would... I think some of these causes we're working on, a lot of them are pretty thin and underdeveloped. But I think a lot of them have potential to become very big causes and very good uses of a lot of money. And I think it's more true of some of these things I'm talking about than it is of the meta stuff.

17:08 Speaker 3: I wanted to turn for a second to ask about... There was a question about outreach earlier. And it seems to me the work you guys are doing has gotten to the point where I wonder how much marginal impact there is to fine tuning the model versus getting more resources allocated to the model that you already have. So... And it's actually a very similar conversation I had at the break, that at this point of the outreach, is you actually spending the dollars right now dollar-for-dollar more impactful by getting additional billionaires behind you versus picking between deworming or something else?

17:44 HK: Sure. So the outreach... The question should we be spending more time on research or outreach. Kind of something we've done from the beginning is we've said, "We want to get the research right first." Because then we have a great product and the we have something to offer people. And then we'll see if people like it and if they don't then we can try and sell it better. But I think the experience with GiveWell top charities has been that, actually, growth has been really strong even without a ton of attention to marketing because I think we did find a lot of people who really wanted this thing that didn't exist, that we created. And we found kind of very early in our existence people were saying, "Your research is good enough. Focus on outreach." And we're saying, "No, it's not."

18:22 HK: And I think that has been worn out. I think that as our research has gotten better, it's been much easier to get people interested. We know what we're talking about more. We have better answers to questions. People give a small amount one year and give a lot more the next year as they realize that they really can feel confident in what we're doing. So that's been our experience to date. We may be biased by it, but I'll also look at the situation for Open Philanthropy and I'll say, "We've not come anywhere close to finding enough good giving opportunities to get at the giving level the venturers eventually want to be at." And as I've mentioned that, the kind of idea of burning down during Carrie and Dustin's lifetime.

18:55 HK: So, I think that idea should be that our focus should be out reach right now. I think for Open Phil is pretty questionable. We do put some low key effort in into just like networking with people who we think are going to be major philanthropists in the past, uh in the future. Yeah, that'd be weird. [laughter] And we have had some progress on that front and certainly Good Ventures is not the only funding partner. We have a funding partnership with Kaitlyn Trigger and Mike Krieger. Mike being co-founder of Instagram. And so, I think that's something we're doing. But I think with Open Philanthropy the challenge of getting a good product, of saying, "We've really thought through what you can do as a major philanthropist. And we've thought harder about it than anyone. And we can share our information." That's a massive challenge. And I think once we get there, especially with major philanthropists who tend to talk to each other and tend to be able to put a lot of thought into their options, I think that we have a good chance of that being all we really need. And so, I think right now just the balance really favors research over advertising or outreach.

19:55 HK: I think... In theory, things may be changing a bit on the GiveWell top charities side 'cause that's becoming a mature product, research wise, and the gaps are getting really big and so you could make a case. If there were no Open Phil, I think we would be definitely ramping up GiveWell's outreach efforts, but the decision we've made is that Open Philanthropy Project is something we want to do and that means a lot of people who used to be really core at GiveWell

have gone over to Open Philanthropy. That means it's a real struggle to maintain GiveWell at the same level. And actually, I think GiveWell's been getting better. But to keep it getting better while people used to be core to it, including myself, are not working on it anymore. So I think, it's like right now priority number one is Open Philanthropy. Do the research. Get the product. Figure out the opportunities. Priority number two is keep GiveWell as good as it is or better and I think it's getting better despite that priority number one and I think outreach definitely is outranked by both of those, and both of those are keeping us busy. One thing I want to add, I just think I just should of said this early on, especially 'cause I didn't go into why we picked the causes we picked.

21:00 HK: I think it's just really important to know that part of the Open Philanthropy vision is that as a major philanthropist, a lot of what you want to do is things that look crazy to most people, but that because you've put in all the thought and because you've put in all the work, you can kind of see the full case in a way that others can't. It's obviously okay if something doesn't look crazy and it's just great, but you should often expect the best giving opportunities that others aren't taking to be things that are a little weird until you've really gotten to know them and think about them. So I want to put out that disclaimer, that I'm well aware a lot of these causes that I'm naming would not be the first thing that would jump into someone's head when they think about the most pressing problems in the world. And that's by design, because we look for things that are both important and neglected. And so, I can't go all the way into how we chose each cause right now, but I will say I'm at least aware that it probably doesn't sound totally intuitive. Kind of might sound like an odd choice of causes, and I do believe that it would take a decent amount of conversation and/or reading to kind of close that distance and maybe, if you're thinking those causes sound odd, maybe you would end up agreeing with us. Maybe you wouldn't after more conversation, but you want to at least flag that.

22:09 Speaker 4: Once a cause is selected from [22:10] \_\_\_\_, what is the medium and long term timeline for [22:15] \_\_\_\_? I'm still a little unclear on what you do after you pick a cause and [22:20] \_\_\_\_.

22:22 HK: Sure. Yeah, what do we do after we pick a cause? It is a little bit case by case, but I think the basic template is, step one, pick a cause. Step two, make a hire. That's the really hard part. Step three, work with the hire. You can call it onboarding or capacity building and get to the point where basically established what the appropriate level of trust and alignment is. And step four is that... We don't know. This is all early for us. Step four, in theory, is that that hire gets more and more autonomy and is able to make a lot of grant recommendations that are really good without needing a ton of oversight or investigation from people like me. Now it may involve that person you need to hire people of their own and I think that's something that I actually think we're still not sure how that's all going to go. That may have to happen. That may not. But I will say, if it were just me looking at all these causes and trying to figure out what to fund, I would not have found or already got a lot more stuff. A lot more giving opportunities than ever could have been found that way and so I think it just continues to be that way.

23:26 HK: And certainly, Chloe... Chloe has a list of things she would like to fund and if she had unlimited money, and it's like she's still getting to know the field in her new role, but it's already at \$40 million and already every single thing on that list is better informed than anything I could come up with for criminal justice reform. So I think this thing starts slow but moves quickly. I think once you have the right person working on the right cause, the money move figures can go up pretty quickly. And the process is that the person recommends grants, we try to ask good questions, but also like depending on how long you've been working with them, there's also a level of trust, and then the recommendation gets made. Yeah, in the back.

24:12 Speaker 5: I'm really [24:13] \_\_\_\_ looking at these longer term [24:16] \_\_\_\_ projection. And two things that occurred to me that you [24:21] \_\_\_\_ when you think about millions of people whose lives [24:27] \_\_\_\_ is marginally more functional. Or another might be ending civil wars [24:36] \_\_\_\_.

24:48 HK: Sure, what do we think about things like improving state capacity in low-income countries, and getting rid of civil wars? I would certainly like to do both. [laughter] I'll say... They're both in a kind of space that I would say we kind of haven't gotten to yet, and so I would say we started with policy, looking at the universe of causes in the US and we started with global catastrophic risks and part of the determinations behind that were practical. Part of it too, is I think if you have a global poverty-reduction lens, I think our top charities are such good giving opportunities that I see a higher priority in pursuing stuff that is the best according to some different lens on the world.

25:25 HK: A lens, for example, that says that some of these other things I've talked about are really the key or really where the most impact is, but I think we do want to get there eventually. I don't have a lot of thoughts at this moment because we haven't looked into them, I don't have a lot of thoughts at this moment on those things. I will say that I... I think they're both things that people have put a fair amount of effort into. I think they're not the most neglected ideas in

the world. And I think effort that's been put in so far, a lot of the results seem kind of lackluster to me, but that's not a very informed opinion and I wouldn't be shocked if we ended up working on some version or some angle of one of those two goals. Sure.

26:03 HK: I don't think they're highest on the list of things I would guess are going to come out looking like they have the best importance, neglected-ness, and tractability, but they're definitely strong possibilities, I'd say. But, it could be a long time before we get to them. We do have our hands full right now with building up the team, the capacity for policy and global catastrophic risks, and then building the step before that for science.

26:27 Speaker 6: So, early... In the first half you mentioned how hard it is to build staff. How time is your most valuable resource that you don't have enough of. And it's clear that you're starting small in Open Phil just because of that problem. But, effective altruism and other related concepts are things that exist outside of GiveWell even if GiveWell maybe doesn't [26:53] \_\_\_\_ everyone else. And of course, part of your openness and transparency is the whole notion, presumably there, is that you want other people leverage what you've done and so, is there... Are you making any effort to sort of be a... Participate in some larger community that's coordinating these kinds of efforts, or encouraging them outside in a way that's not just hiring staff internally?

27:23 HK: Sure, are we engaging... One way to put this question is, how much are we trying to engage in community building as opposed to capacity building, right? Building our own staff versus kind of trying to build the community of people who share a lot of the similar values. Again, I think focus is very important. I think having a contained set of goals is very important, and I think right now we've chosen our goals. Which is like, we have this idea we're very excited about which is building a framework to figure out what a philanthropist should do in policy or global catastrophic risk, something non-meta.

27:54 HK: Something that's not about increasing the supply of people who are like us, but something that's about instead getting to the bottom of a problem and being able to share our reasoning and create a model for that. I think that's our best comparative advantage. I think that's really what we're best at, is taking these kind of overwhelming ocean-boiling type problems with a lot of confusing information and evidence and figuring out a way to cut through it, come out with recommendations we can justify and then help people make decisions. That's the priority, that's what we're doing.

28:22 HK: Effective altruism is a community of people that I think there's quite a bit of overlap between our values and theirs. There's a lot of interest in this kind of idea of doing the most good possible, and there's also a lot of stuff that goes on in that community that we don't agree with, or aren't necessarily comfortable with. So, there's a lot to be excited about and I think we're generally positive about the community and do participate in things.

28:46 HK: I think we interact with the community, we network, we go to events, but it's... I think to say we want to actively grow that community would turn us into a community-grower instead of what we're trying to be, which is a developer of knowledge and arguments about philanthropy, and that's just not the path we've chosen because I don't think that's what we would do best. If you're interested in more, there's a post called something like "Why the open philanthropy project is not funding groups focused on effective altruism." And it's on the effective altruism forum, it's not on our blog, but if you Google it, you'll find it. It goes into a lot of detail.

29:23 Speaker 8: Do you find anything useful from anybody else's efforts that they've stayed transparent?

29:28 HK: Yeah. Do we find things useful from others' efforts? Absolutely. It depends on exactly what we're talking about. If you're asking about the effective altruist community, I would say absolutely. There's a community of people who are really intense about doing the most good possible, and they raise all kinds of arguments that I wouldn't have thought of myself or that the GiveWell staff wouldn't have thought of just all internally, and we really appreciate that and we're generally friends with people in those communities, in that community, talk to them a lot, get a lot of ideas from them, and then try to think about which of those ideas make the cut as the most promising ones when they're up against all the other ideas in the world. So yeah, I think we've got a ton of value that way. I mean if you're asking the broader question, I don't know, we really rely on the rest of the world to produce a ton of useful information, and we're always trying to do the least work ourselves possible, and just take what's already out there and make sense of it. That answer your question? Other questions? Yep.

30:27 S2: Just wondering from previous question if you'd done much work on sort of immigration reform, if you're sort of just focusing on US Policy, and that being an alternative to state reform in other countries?

30:41 HK: Sure, immigration reform, yeah. Immigration reform is like one of our focus areas, and we think it's a great area because the importance is really off the charts by US Policy standards. So basically the benefits... Most immigration debates tend to focus on, we have a bunch of people who are here and undocumented, what's the path to naturalization? They don't tend to focus on how many people get to come across the border. We think in some ways that's the much more important question. When someone comes across the border, same person, same skills, but they're often going to earn 3X as much, 10X as much, sending money home to their family. You can imagine that an increase in immigration, that would be kind of not huge in the scheme of things, could really just have a massive impact on the well-being of a lot of people. And so we see it as an area where it is politically very difficult. I mean it's tough. I mean it's a federal issue. You're not going to be able to pick your battles at the state and local level, you gotta do it basically fight your battles in DC, and that's not a great place to be trying to get stuff done right now.

31:42 HK: So it's politically very difficult, but we think it's worth it, because in the long run, who knows what'll happen? And we think there's at least enough of a chance, and we think the possible benefits of really kind of immigration friendly immigration reform would be huge. But it's surprisingly not a huge space. So I think if you're... If we wanted to open up a program about naturalization and path to citizenship, there'd be tons of organizations to fund, there'd be tons of people to hire. If we want to focus our programming on crossing the border, especially pro-poor migration, so people coming from poor countries to rich countries, we really looked for people to fund, and have had trouble finding them, and have no idea who we would hire, or frankly what that person would do all day. So that is one of the three areas that Alexander works on. And it's... That way we're able to keep our eyes open.

32:38 Speaker 10: In the category of [32:39] \_\_\_\_, I'm curious what you think about climate change, or how you think of it as a risk that clearly has multi 10s of percent chance of being a very bad outcome. Why [32:54] \_\_\_\_.

32:55 HK: Sure, so where's climate change on the global catastrophic risk? So I think if you look at climate change as a whole issue, I would say that the importance is very high. I mean you can't name a lot of single phenomena that could have kind of a double digit percentage impact on World GDP down the line. And I think the models are not reliable, and that, to me, makes it scarier, 'cause who knows what's going to happen. I think climate change is on a fairly short list of things that could be most disruptive to civilization. I don't think it's higher on the list than the two that we have prioritized. I think pandemics are more threatening in the short and long term in terms of the amount of damage you can imagine, kind of within some realm of some kind of probability, some kind of confidence interval.

33:40 HK: I think AI is a much harder one to really get your head around, but I think the extreme cases for that are worse than the extreme cases for climate change, and also could be on a shorter timeline frankly. So, I think climate change is very much up there on importance, but I certainly don't think it's clearly more important than the other two. I think it may be even like not quite at that quite same level, though definitely high. And then the neglected-ness I think is way lower. That is one of philanthropy's favorite causes, it getting hundreds of millions of dollars a year devoted to it, and I don't know really what is left to do, in terms of philanthropy. I mean there's a ton left to do in terms of policy, but what more can we do to try and get that policy to happen?

34:25 HK: So, that's climate change as a whole. I mean I think there's this sub-category of climate change around solar radiation management, which is kind of this idea of a worst case response, where you would try to do these kind of highly risky things that might cool the planet, and you can imagine it being the thing that prevents the harms of climate change, you could also imagine it being the thing that makes everything much worse, and is actually the danger itself. And we think that researching how to do that safely, and how to govern it, so that you don't have a war breaking out because one country decided they want the whole planet to cool down, is actually a really under attended to area. It's actually pretty important, and it's very neglected. And so we have made some grants there, and I would call that basically third on the list of priorities, but we have capacity for two right now.

35:19 Speaker 11: I don't know if it's too specific a question, but what are you doing about pandemics? What is your focus, and what's your approach to the issue of pandemics?

35:27 HK: Sure, what's our approach to the issue of pandemics? I mean our approach with a fairly big and robust cause like that is basically to try and outsource it to someone awesome that we hire. And the way that we try to do that is we try and get a really broad picture of the space, and understand all the things you could believe, all the things you could be excited about, and then we interview people, and we want to hear from them, we want to get a sense from them, that they kind of have thought with the same breath as us, but much more depth, and more intelligence, and feel like they can make those trade-offs better than we can. I would say that's kind of the rubric for how we ran the criminal justice



reform search, and that is how we are running the bio-security search. So we're not here saying, "We figured it out, we need to get DNA sequencers installed everywhere in the developing world." Or whatever it is. We're saying we have a sense of the possibility space, and we're looking for a person who convinces us that they're thoughtfully trading things off.

36:25 HK: But one of the things that we have come to believe, is that it is a very government-dominated space right now. There is very little non-profit philanthropy presence. And I think we're starting to feel some of the distortions that causes where just people are... It feels like a lot of the people we talk to are by default, focused on things that seem politically important, like the threat of Ebola and things like that. Rather than why do we... If you think about the very worse things that could happen, and you ask whether their probability is correspondingly lower enough to cancel that out. We think the answer is no, and we've had trouble really getting, finding people who we feel are appropriately knowledgeable about and thoughtful about that particular aspect of the problem. So, you're making a face, so if you have ideas... Yeah, refer us.

37:17 S2: Well, I'm just thinking, you know, it's pretty simple, pneumonically spread Ebola.

37:21 HK: What?

37:23 S2: Pneumonically spread Ebola.

37:24 HK: Okay, can you elaborate?

37:26 S2: Yeah, I mean Ebola was transmitted through body fluids, like AIDS. Let's just say that that had mutated and it's a possibility, a very good possibility.

37:40 HK: Oh sure.

37:42 S2: And it mutated to a Pneumonic spread, in which all you have to do is breathe.

37:46 HK: Sure.

37:47 S2: You know, how do you deal with that? I think the World Health Organization is a basket case. This is something that could happen tomorrow, and nobody's prepared for it.

37:57 Speaker 12: A deadly flu is the other one, if you think about it. The other way, a deadly flu...

37:57 HK: Say again.

37:58 S2: A deadly flu would be [38:00] \_\_\_\_\_.

38:01 HK: Yeah, I mean I think yeah.

38:03 S2: A pneumonically spread fatal disease is very possible.

38:07 HK: Sure.

38:09 S2: Very possible tomorrow, and nobody is really... I mean this is...

38:14 HK: Right. So that I agree with, and I think that our focus has been more on flu, than the idea of Ebola mutating. You could have, yeah, there's certain varieties of flu...

38:23 S2: I'm just talking about fatal pneumonic disease.

38:25 HK: Right, right. And there's precedent for massive death tolls from that, and there's possibility, according to models, of something even worse. So I totally agree, and then like, finding the person who we feel like is going to be as focused on the worst version of that, as they are on the politically salient stuff, that's been tough. So it's...

38:44 S2: Well, influenza was not a very... It did have a very high mortality rate. Just figure that mutation but double

the mortality rate of that. Bam!

38:57 HK: Bam! [chuckle] Yeah, no, I agree. I mean, that is one of the main things that we think is a good example of a worst case, though it depends on the exact specifics. I mean there are flu pandemics that are more likely less scary, and getting reasonable amounts of thought, and then there's like versions of them that we think are more scary and less likely getting less thought. Anyway, I think I should move on to the next topic, and also probably wrap this up soon, but I'll take one or two more.

39:24 Speaker 13: Do you guys think about finding and hiring, like the program officers, and then building that trust, because it's like you said, there are areas where you don't know that much about it. That's why you're hiring them, and it seems like it's really important to get right, really hard. What's your guys...

39:36 HK: Yeah, how do we think about evaluating program officers, given that we don't know the area as well as they do? I mean, so the approach has been just to basically ask a ton of questions, and kind of go a little bit overboard with the skepticism and the poking early on, and see what happens. And so, what we want to see, is someone who when they first bring us a proposal for a grant, it's like they've already thought about all of the stuff we're going to ask. And the more we get... And obviously that's not going to happen the first time, right?

40:07 HK: But, the more that we kind of build things up to that way, and this is something that I think, in a certain way, we have a lot of experience with 'cause the GiveWell top charity process is like this deep dive, that is deeper than most people attack most questions, and most cases. And we'll go ridiculously deep on these set of charities, so that we can really stand behind them, and so we are fairly used to this idea of like, "Hey, go check out this charity, come back to me, tell me how good they are." Then I'll ask you like a million questions, and then after we do this like a hundred times, you'll learn you've already asked the questions, and then by the time I get to you, you've already thought it through. And there comes a certain point in which you're like, "Hey, I found this charity. They're good." And I'll be like, "Yeah, okay then they are." And so that's, it's a similar process we're trying to follow, yeah, and I mean, we're not experts on the areas these people work on, but we certainly know how to ask a lot of questions. So that's one way to get at that.

41:04 HK: I mean you said this before, but we learn a lot about an area at a high level...

41:06 HK: That's right.

41:07 HK: Before beginning a search, so that we know which questions to ask.

41:13 HK: Yeah, as an aside, one of the common pieces of advice we get from Philanthropists, is to worry less about the project and the issue, and worry more about just like the quality of the people. But, the other thing is that we feel, at least from our hiring, or how we've at least thought about how to do it, it's important to have a certain level of subject matter knowledge to evaluate the people. And so that is definitely a part of it, and for us, we feel like we have to get up to speed a certain amount before we can even engage in these conversations. And yeah, that's definitely a part of what we do. I can take like, let's say, two more, yeah.

41:53 S 14: I'm wondering about, sort of, politically loaded topics, where you can imagine that you donate a million to pay for this bill, and some other group thinks this bill is going to do all these things, and they donate a million to fight against this and you would have been much better off saving your money than spending \$2 million on the big charity. And is this something where you have enough confidence in your convictions than like I probably got this from a helping those people standpoint, that you're okay taking that stand or is it something where you really try and shy away from?

42:19 HK: Sure. Do we try and shy away from things when there's someone on the other side? Certainly you can pose this hypothetical and you can say, well, instead of you fighting for X and someone else fighting for not-X, why don't you both get out of there and then you get the same affect. And then you all donate to charity and the world's better off. I think one problem with this is that generally when we look at policy issues, we are trying to pick issues where the non-humanitarian side has all these natural reasons to exist and the humanitarian side doesn't. And it won't if some philanthropist doesn't make it. And so I think that really does remove quite a bit of the symmetry. And I think that's usually what we are looking at.

42:57 HK: And so, for an example, with Criminal Justice Reform, a lot of what you're dealing with is prosecutors as an

interest group. You're dealing with the status quo. I think you're dealing with a fair amount of uninformed public opinion. Those are things that you're not going to make a deal with them. Just give to charity instead of doing this. And really, what you've got is, one side of the issue is standing in the way for reasons that have nothing to do with philanthropy or humanitarianism, and the other side will only exist if some philanthropist comes in and causes it to happen. So, I don't see it as being fraught in the same way that some people imagine it. I think usually we're coming in on issues where the balance of power is naturally lopsided and we need to kind of like unlopsided it. We're not usually coming in on issues where our side is already winning and we just want to pile on until it's some overwhelming resource superiority. I don't really see it that way. Certainly, I think... I'm less certain about my views on any given US policy issue than I am about the impact of deworming, and as you know I am completely uncertain about the impact of deworming.

[laughter]

44:06 HK: So there does have to be a discount and that is something we have to keep in mind, but also, I think policy can be very leveraged. I think if you look at historically, some of the best things philanthropy has done and just some of the best things that have happened, period, a lot of them look like people getting up to fight for a change. In public policy, or just fight for a change. I don't think I want to leave that stuff on the table just because we're uncertain about it. So I think that's what open philanthropy is for. It's to do that stuff that we are not certain about it, but we think it is worth trying anyway. How it compares the top charities, I would say that's an open question. Last one.

44:44 HK5: Just wondering about the criminal justice reform. I've seen in the past a GiveWell talking about [44:50] \_\_\_\_ I'm wondering what kind of metrics you ultimately wind up [45:10] \_\_\_\_ . How much of that... Where right now as you're moving in large trends [45:22] \_\_\_\_ . Where do you see that relative to say, [45:26] \_\_\_\_ . If right now, you say that [45:36] \_\_\_\_ , at what point would you think you would you get the knowledge to say in retrospect, this will be better or worse.

45:45 HK: Okay, so we're going out with a bang because that was about three questions. There was a question about the dallies. What is the metric that we measure policy grants by? I'll give that answer first. You can talk about the impact of policy change on peoples' lives. You can describe it in terms of some kind of version of disability, adjusted life years. So, a calculation we have done in the past is we've said a year in prison is like 50 percent as bad as losing a year of your life or something like that. And so you can do those calculations to compare new or top charities. Or you can do it in terms of dollars. So you can try and convert those into dollars. There's methods for doing that. And you can also look at the money saved by governments when they're not spending as much on prison. So, I don't think we've really settled on one or the other.

46:34 HK: I think we just tend to always, when we ask a question, we look for whatever units we can answer the question in. I think something we've been tending to do more with policy stuff is calculate kind of like dollars of value added to society divided by dollars we spent. And we call that the return. And then that's something that's very messy and very imperfect, but at least lets you get ballpark comparisons. So with criminal justice reform, the early stuff we're looking at is like a few hundred X return according to our estimate. And then cash is like, give directly, I had mentioned before. You could think of it as a hundred X return. And so that would kind of imply if you just took it at face value that the criminal justice reform stuff is better than cash transfers that may be only even with or not as good as bednets and deworming. I still think we have a lot to learn here. And I think that idea could go up or down a lot and so we're not being real hard core about that yet. And that's what I said before about this idea that we hold ourselves to a lower standard when we're getting to know a field. So what is that going to change? I think what we're hoping for is a world where we gain confidence that our working Criminal Justice Reform is best in breed or as best as we can make it. And we found the best giving opportunities we can. And we've understood as much as we can about how good they are and then we'll have the confidence to make that.

48:00 HK: Adjustment and you know that's an evolution that's happened before so with deworming, you know, we have gone from kind of feeling like we had a crazy spreadsheet that no one really gets to feeling like we've thought about this enough and understand it enough that we can start to really believe some of these numbers. To a certain extent, I think we're going to go through the same evolution on criminal justice reform. I think it'll take a while and we have to plan for that fact but I do think it'll happen. I mean, I think certainly like five years from now I'll be very surprised if we don't have like a pretty good opinion on how those things compare and then... I mean, I think, look, there's a lot of judgment calls here and so we may decide it's ambiguous, it's in the range, and in that case, we are looking for places to put a lot of philanthropic money and different people have different interests, and so I think if it's ambiguous you kind of want to maintain the program.

48:44 HK: If criminal justice reform looks better, you probably still want to keep to the status quo because I think there's just a lot of, you know, there's a lot of people who you're not going to get to give there over top charities, and if criminal justice reform looks kind of like robustly worse or a lot worse, then we would look to close the program. We would transition out. I mean it wouldn't be overnight, but we would transition out. So with that, I am going to call it but I will stay around and, you know, if people want to come up and ask questions I'll do that for a little bit. So, thanks everyone.

[applause]