

GiveWell San Francisco Research Event, December 15, 2014 - 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.

00:14 Holden Karnofsky: We're going to get started again. We're running a little late, so if you guys could quiet down soon. Like the next 10 minutes, 20. Okay. So, yeah, we're getting started again. Sorry, we're running a little behind tonight because of the pizza, but I'll just talk fast. So, I'm going to talk about the Open Philanthropy Project. This is our newer project. It's exploratory. It's been going for a couple years. Definitely less well developed than GiveWell's top charities, and I think in some ways it's analogous to our GiveWell top charities were in 2007, 2009. Basically, the real quick rundown on Open Philanthropy Project, 'cause there is certainly a lot to talk about, and I don't want to spend too much time on the basics, but we... A few years ago we ran into Good Ventures represented here by Carey, the president. A major foundation funded by Dustin Moskovitz, one of the co-founders of Facebook. And just heard the problems they're thinking about as they're trying to give away a very large amount of money over the course of their lives. And thought, "This is very related to the questions we think about, and it's the kind of thing we like to think about, "How do you give away money as well as possible? And is there a way to do that in an open and transparent way so that others can learn from what you're doing?"

01:27 HK: In some ways, it's very continuous with GiveWell's top charities because both projects are about saying; "How do I give as well as possible, and how do I share that knowledge?" And both projects are a reaction to asking that question and not finding other great sources of information, which I think both we and Carey feel that that has been our experience. Big difference is just the kind of giving opportunities we're looking at are completely different, and the criteria and the process we're using are completely different. So, GiveWell looks for things where we can fund something that's pretty direct, something that's supported by evidence, something that has worked before. Open Philanthropy Project is more open-ended. And kind of as I mentioned in the last session, a lot of the things that we talk about for Open Philanthropy, frankly, it's going to be hard for me to give a full sense in this meeting of why it is that we think a particular area is promising, because a lot of it is built up on frameworks and informal knowledge that is just tougher to put onto a website and tougher to absorb, although we're doing the best that we can.

02:32 HK: So, what we're not doing is we're not looking for things that are supported by randomized controlled trials or things where we can precisely quantify the impact. What we are still trying to do is be systematic. And we followed a basic framework in which we, A: We've decided that what we want to do is start the project by thinking about what areas we want to work on. We call them causes. "What problems do we want to work on?" We think about that instead of thinking about what organizations we want to support, because I think when you're doing major philanthropy or open-ended philanthropy, one of the things you can do is create new organizations, or change the way existing organizations function, or just go out there and be known as a funder, and see new giving opportunities emerge that you couldn't have found out about otherwise.

03:17 HK: Furthermore, just we've generally found that when you get away from randomized controlled trials, even looking at one project or one organization, it's like you'd have to get to know a whole field and a whole ecosystem in order to evaluate it. And by the time you do that, there's a lot of other organizations and projects you'll be able to understand. So, our starting framework is to say, "What causes should we work on?" And the general operating principle, although the specifics

can vary, is that we generally will take a very large list of causes. A large list of potential problems in the world we could work on and try to find the ones that have a good combination of being important; meaning they affect a lot of people and affect them a lot; being tractable, meaning there's something we can do about the problem; and being uncrowded, meaning there isn't too much attention to the problem already. And so, our ideal cause is something that doesn't get enough attention for how important it is and for how tractable it is.

04:13 HK: A little more specifically... And then, we do this in a hierarchical way where we first look at, sort of similar to the top charities process. First we look at a cause at a very low level of depth. We might have a few conversations, read a few papers. Then, we take the promising ones and look into them in more depth, maybe talking to 10 to 30 people. And then, when we get really interested in a cause, we start going around and trying to make some grants and getting a feel for what it would really look like to work on that problem and what kind of giving opportunities you would find. We've divided the philanthropic world up into four areas that cover a lot of the causes we'd want to work on. Not all of them, but a lot of them. So, one of the things that we think about a lot, one category, is policy. How can philanthropy influence policy? If you can... And I think one of the big theories of how philanthropy can have a big impact is that you can spend a relatively small amount of money to help organizations get an issue on the map, or change minds, or come up with ideas for how policy can be changed. And then, when policy changes, especially at the federal level but even at the state level, that can be an enormous amount of money and an enormous amount of impact relative to what you can directly pay for.

05:29 HK: We are focused on the US policy right now. A lot of that, most of that is just practical. It's just because I think to work in policy, networks are a ton of what you need, and I think it's just much more practical to build those networks and have the necessary background knowledge in the US, which is where we live, and the ecosystem we understand best. I also think the US is... I can't think of a clearly stronger candidate to work in. I think there are some countries that would be just as good to work in, but none that are clearly better, when I just kind of step back and think if I were starting from scratch and didn't care about where I lived.

06:05 HK: So, there is US policy and there's a whole bunch of different policy areas we can work on, and I'll talk about that. But another category is global catastrophic risks. This is a good opportunity for philanthropy, because there are some problems in the world, climate change probably being the most well know, but also pandemic risk, asteroid risk, a whole bunch of other stuff, where there's kind of a very low probability of something that can disrupt sort of the whole global civilization. With these kinds of issues, there's not necessarily a particular company or sector or even government who's job it clearly is to worry about these things, and so in some ways, we think they're a really good fit for philanthropy. And again, if you can have a small impact on a chance of a really large, massive, global threat, that can be a way to magnify the impact of your money, so that's a category.

06:56 HK: Other categories, scientific research. I think when you look at the list of, and one of the things we put a fair amount of work into is understanding what philanthropy has succeeded at in the past. Understanding the history of philanthropy, you see a lot of cases in which people funded the creation of new ideas and then those new ideas can be used without limits. So that's another way to magnify the impact of your dollars, and I think it's been a fruitful area for philanthropy in the past.

07:20 HK: And finally, global health and development. Obviously, we're very interested in global

poverty. Our top charity's work is very relevant to it. But we also wonder, what can a philanthropist do about global poverty besides doing very direct delivery of interventions? One of the examples of this, we are really into bed net delivery, but as we've studied up on it and as we've learned about it, we've kind of wondered why isn't there better data in general to help guide the delivery of bed nets? Data such as who's using the nets and could that be a leveraged opportunity to make a difference? So that's the kind of thing that we think about.

07:54 HK: First I'm going to talk about our work in US policy, and I think that's going to be the easiest way to understand the general framework that we use. And then I'm going to take a break. And then a bit later, I might talk about some of the other categories, depending on how much time we have.

08:11 HK: US policy, just in terms of the process. I mean, we spent a while just kind of trying to get basic feel for what a philanthropist can do to influence policy. How the whole system works and how we should think about what makes a good issue area to work on. And what we've kind of ended up doing is we've canvassed a broad list of policy areas and learned a little bit about a lot of them and most of this work has been led by Alexander Berger, who is here. And we've kind of looked for issue areas, that again they're combining. They're going to... We're looking for things that are outstanding either in terms of how important the policy area is like it affects a lot of people or whether there are unusual political opportunities to work on an area or like how crowded the area is. And so, all of us, equally, would rather work on policy issues that are getting less attention.

09:03 HK: One of the things that I want to put up front before I go into the list of causes, is I think, generally, we have been counseled by experts. And I think that we also believe this is right. To be a little bit less obsessed with political opportunity than you might think one should be, I think, coming in. I think a lot of people when they are first thinking of the idea of political involvement, they think where are the opportunities? Where are the things where I can actually get something done? Don't want to do something that's too politically tough. And I think what we've kind of come around on, is that a lot of what philanthropy is going to do best in this area, is long term stuff like building advocacy organizations or building ideas that will then be ready to go when the political moment comes. And it's hard to predict when a political moments going to come. So there's just many cases of... There's a point in time in which people say a certain policy issue is impossible. And then, no, it's not. The landscape can change quite dramatically and quite unpredictably. And what really matters is that you got an infrastructure there and people there and ideas there ready to take advantage of it.

10:09 HK: The issue areas that we work on... So, I'll give examples of some of the ones we're most interested in. Some of the ones that have stood out to us. One of them is criminal justice reform. Basically, the short story here is that the US has a very high incarceration rates; historically and globally, has like relatively low crime compared to what it's been in the past for the US. And there's these questions of whether we can reduce incarceration, which has incredible humanitarian costs as well as fiscal costs, without incurring too much public safety cost. The thing that has made this cost stand out is when, when we've gone canvassing people, this is the one that has stood out where people have been saying there is political opportunity here; and it's real, durable political opportunity. It's not the kind where there's a headline, an issue's hot for three months. By the time you decide to work on it and fund some organizations, it's gone.

11:02 HK: This is the kind of thing where I think things have changed in the criminal justice world.

State budgets are tight. There are major conservative voices now. People on the right not just on the left saying that we need to reduce incarceration. That it's big government. That it's expensive. That there are better ways to keep public safety. And so, this is an area where I think, this is the one that I would point to as the biggest window of opportunity, where kind of robustly, over the next several years, maybe beyond that, I think there will be unusual opportunities to get something done.

11:33 HK: Other causes that we are interested in. Labor mobility is in some ways the opposite. This is this idea that the best anti-poverty intervention we know of, is when someone steps across the border from a poor country into a rich country. And just by virtue of stepping across that border, their job opportunities completely change. The labor market within which they are operating completely changes. And their life opportunities, and their abilities to earn, and their income change a lot for the better. And so this is, unlike criminal justice reform, don't see a ton of political opportunity here. Immigration is a pretty polarized issue. It's been pretty frustrating, the lack of progress on it. However, also unlike criminal justice reform, I think the magnitude here is enormous. The benefits to someone stepping across the border are really huge, and even a small change in immigration policy could have really big benefits by our calculations in humanitarian terms.

12:24 HK: Also, immigration policy is generally a space that gets a lot of attention, but the specific issue of the benefits to people coming from poor countries into rich countries, especially for low-skilled, relatively low paying jobs, I think is an area that gets a little bit less attention. And I think there are some places in which we found things to fund that I think wouldn't have gotten funded otherwise and are pretty basic attempts to make that happen and to focus in on that angle of immigration policy.

12:54 HK: Another cause that we're very interested in, and this one is a little bit of a mouthful, "Macro Economics Stabilization Policy." The basic idea here is that there's a lot of debate over what we should do when there's a recession, like a prolonged spell of poor economic growth, and unemployment, and there's certainly a lot of debate about it. But I think one of the things that we have come to believe is that a lot of times policy and especially the Federal Reserve, is trying to trade off against risks of inflation versus risks of unemployment.

13:28 HK: In humanitarian terms, from a humanitarian perspective, we think that unemployment is often a bigger threat than it's given credit for by a lot of the decision makers. This is another case of an issue where the importance is so huge that relatively small impacts can be very big. Recessions are always national and often global and a small impact on one could really have a huge humanitarian result. And this is another issue where... In some ways this issue is crowded.

13:57 HK: It's a very, very popular topic among economists, among economic bloggers, for sure. There's an army of intellectuals and analysts at the Federal Reserve. In many ways, the more political aspects of this issue, I think, get a lot less attention. The idea of real advocacy on behalf of people threatened by unemployment, the idea of creating policy proposals that'll be ready to go next time we're having a debate over stimulus, the idea of trying to make our fiscal policy, in other words, how much the government spends, trying to make that do a better job responding to recessions without necessarily allying ourselves with wanting bigger or smaller government. I think these are things that we think there may be some space. There may be some lack of crowdedness, and we think that's pretty interesting.

14:45 HK: Quickly name some other ones... There's land use reform also known as zoning. I think is an example of an issue that's neither the most important issue in the world, nor the most tractable, but it's got so few people working on it that it's interesting to us, and there's several issues where we kind of look at it, and we say, "Wow. It would be really interesting to take a policy space where there's almost nothing going on, and get something going on," because in terms of advocacy, I can imagine the returns to that being especially big and especially good.

15:14 HK: We remain interested in foreign aid policy, but also perceive that there's a pretty good infrastructure in place already. There's pretty good organizations that advocate for more and better foreign aids. So finding what we would bring to the table, what we'd be able to add with more money is a little bit non-straightforward, although we think about it. We think about treatment of animals in industrial agriculture and more generally just the problems caused by factory farming as another issue that we think the importance depends a lot on philosophical judgment calls. But it does get a lot less attention than we would've guessed when we first started looking into it. It's not the world's most partisan or intractable issue.

15:53 HK: One other factor that I haven't mentioned is... I think in a lot of cases, and especially in today's environment, if you can work at the state and local level, that's often an advantage. So if there's many, many places you could go to try and get a win, and you can pick your battles, that's great. If you're kind of "go to DC and get a win there or go home," I think all else equal, that makes an issue more difficult to work in, although some of our issues are definitely DC only. So that's like a very broad overview. I think I can take any of these causes and probably we could do a whole meeting on it, and in some cases we have. I will take questions now and then maybe move on to global catastrophic risks after some of that. Yep.

16:32 Speaker 2: So it seems like with the exception of zoning, all of these skew left. Why?

16:37 HK: With the exception of zoning all of these skew left. Why? I think we are definitely skewing left on more causes than we're skewing right on causes. Although it's certainly not a hundred percent...

16:50 S2: Migration.

16:51 HK: I think some of the immigration stuff is... It's not really clear left or right. One of the groups that there was an early grant to from the Open Philanthropy Project is definitely a group that's on the right, Immigration Works. I don't think it's 100% accurate. On the other hand, we come in to US Policy with no interest in being bipartisan or non-partisan. Our interest is in accomplishing a lot of good, and in doing a great deal of humanitarian impact. There's no particular reason to expect that that's going to leave us evenly split between the two parties. The parties represent different interests, different coalitions. I generally think that in the US, the way the US political spectrum works, the left is just a more natural home for utilitarianism, for example.

17:39 HK: I think the right often has values that is going for terminal goals that are inconsistent with utilitarianism. And often the left, at least in rhetoric or in some of the values, is more consistent with that. I think that we have a set of values and a set of policies that does not fit neatly into any one existing platform. I don't think you can call it progressive. I don't think you can call it libertarian or definitely can't call it conservative, but I'm also comfortable with the fact that it's not equally divided between those things. It's closest to progressive, and I think that's because the

values are closest to progressive values. I think we care more about helping people who need help and have disadvantages. That's something we care about a great deal, more about that than about the abstract concept of liberty or about the values of conservatism. Yeah, Jacob?

18:31 Speaker 3: So for labor market mobility, it seems like naively if you allow unskilled, low-paid labour into a country that would increase inequality among people already living in that country because low-income people would have their wages driven further down, so how do you take that into account in your analysis?

18:53 HK: Sure. So how do we think about the impact of immigration on inequality and in particular this idea that if you have low-skilled immigration into a country, it could lower the wages of other low-skilled people via competition and then lead to greater inequality. So a couple answers to this. One, I think we keep our tendencies from GiveWell top charities of being very curious people, wanting to learn as much as we can about every question, constantly criticizing ourselves, reflecting on what we're thinking, and trying to find areas where we could learn something might be wrong.

19:26 HK: So, with immigration, this is one of the questions that we commissioned a pretty good research project into. David Roodman did a write up on it. I think if you Google David Roodman, Immigration, GiveWell, you'll probably find it, where he reviewed the literature on the impact of immigration on native's wages. And he concluded that, and I think he did a good job, that there is not that much reason to expect a really big negative impact on wages from immigration. Immigrants, yeah, they provide competition for jobs, but when they get income, they're also providing demand, and in the long run, basically, you're just growing the size of the economy.

20:04 HK: And then as far as impact on low-skilled people particularly... Yeah, I think the effects that are there in the literature don't look huge even in the short run, and in many cases, it's more of an effect on existing migrants than it is on low-skilled Americans as a whole. But the other thing I'd say is that we are generally more disposed toward well-being than toward equality as an end in itself. A lot of why we are interested in inequality is because reducing inequality can be an efficient way to increase total well-being in some sense. And I think that you can concede some increase in inequality when people step across the border, but you also have to look at people going from being extremely poor to being much less poor by global standards, and I think that's just a huge benefit. We don't want to leave that on the table. I think a lot of existing political ideologies tend to leave that on the table because they tend to look at what the well-being of Americans looks like instead of what the well-being of the globe looks like. That's not the way we want to go.

21:07 HK: And obviously, early in this session, I'm already talking about our values, what we believe, what we care about a lot more than we were with GiveWell top charities. There's a lot of value judgments with both. Philosophy is important to both, but it's certainly the case that as we get into Open Philanthropy Project, we have to make a lot more of these judgment calls and a lot of what we're doing has to do with what we believe in, what we care about, and that is why we're giving it a different name. Eventually looking to separate the organizations, make them totally different projects. What they will have in common is this drive to do as much good as we can and a drive to talk about it and share what we're learning. Yeah?

21:42 Speaker 4: So let's talk about the... Another difference. Where is the Open Philanthropy Project in terms of advising individual donors on how they can deploy their money?

21:53 HK: Sure.

21:53 S4: Like I know you're at the stage where you're giving things away, and that's partially a way to get in better conversations and stuff like that, but...

22:00 HK: Right. So where is Open Philanthropy Project in the process of getting recommendations to individual donors? I think we're not there yet. Right now, open philanthropy project is a collaboration between GiveWell and Good Ventures, and the grants will come from Good Ventures and we have a process for deciding on them where the ultimate call is Good Ventures, but that's taking into account advice from the GiveWell staff.

22:29 HK: Good Ventures is going to have a good amount of money to give away, and I think we are trying to build a very robust and well-scaled operation for giving a lot of money away really well, and so the thing we're not trying to do is as fast as possible get some recommendations out for moving \$2,000,000. The thing we are trying to do is take our time getting to an operation that can move billions of dollars. So that's our focus and I think we're a ways away.

22:57 HK: If someone came to me for a specific cause or for general recommendations and said, "What are organizations that you supported or seen that I could consider supporting?" I would have some answers, but they would be tentative answers, and it's not our main focus. I think we will get there, and I think we... The more that we've got this operation developed and the more we know what the giving opportunities are, we'll think about how can we bring individual donors into this, but I just think that's a little bit in the future for now. In the back all the way.

23:27 Speaker 5: Could you give us a sense... I run a shop that makes investments, both debt and equity and some grant making in the policy space. Could you give us a sense of like the kinds of organizations that you can imagine giving to? Are they all think-tanky or are they campaign-oriented or something else entirely?

23:47 HK: Sure. Can we give a sense of the kinds of organizations we're interested in giving to? Are they think-tanky? Are they campaigny? Are they something else? We think one of our advantages is just the flexibility and the ability to look at kind of everything, and a lot of what I have tried to do is get an understanding of like what does the complete toolbox look like, right? What are all the things you could do to try and change policy? You could fund people to run campaigns. You could fund people to do grassroots lobbying, which means mobilizing large numbers of individuals. Or you could fund people to do sort of, go around on in Washington and lobby individuals. Or you could fund think tanks to develop policy proposals that are kind of have some intellectual thought behind them, but they're also designed with policy and politics in mind on like academic work. You can also fund academic work to try and get change people's minds on issues. And there's more on the list.

24:42 HK: And a lot of what we do is we'll looking at a cause, and we'll say, "This cause is really important. It has A, B and C, but it doesn't have D and E. That's pretty interesting." Macroeconomic stabilization's an example that. I mean, there's a ton of academic research. There's a good amount of various kinds of things, but in terms of grassroots, especially the side that we are inclined toward, there's like almost nothing. And even like think tank work, certain kinds of think tank work, I think there could be more than there is. I'll try and go to the side this time, 'cause I keep just looking

ahead of me.

25:15 Speaker 6: There's two spaces I'm wondering if you've looked at. Have you looked at all of that positive psychology, happiness, well being that sort of change? And have you looked at using the media for inspirational messaging to inspire more people to be involved to care about the...

25:32 HK: So have we looked at media to inspire people to be more involved, care more about these issues? Have we looked at positive psychology, science of well being? On media, we have talked a lot about the general tool of media. I think of it more as kind of a tool than as a cause. So, kind of the way I was just describing it, there's like causes are issue areas that you want to get changed. And tools are ways that you might change them. I think of media more in the second category. And so, what I've been trying to do is just remember when we're looking causes that I think one of the interesting things to think about is, can you fund media that's going to make a positive contribution of things? And the example of criminal justice reform, there's a relatively new group, called the Marshal Project that I think is going to be trying to raise the game of the journalism going on around criminal justice. And I think that's an interesting opportunity. Something we haven't supported yet, but we're keeping an eye on and we could. Positive Psychology, I'm not sure exactly what you're referring to. I think what you're referring to sounds more to me like research. It sounds to me more like funding research into understanding the determinance of well-being in positive Psychology. No? So, what is it?

26:42 S6: I was thinking there are really good interventions for personal change and for well-being that are really under-utilized.

26:49 HK: Interventions for personal change and well-being that are under utilized. So is this a policy issue? Like a...

26:56 S6: Like a...

26:59 HK: Like getting people to do things.

[overlapping conversation]

27:00 S6: That could be funded better.

27:02 HK: Yeah. I do think this is an interesting area. The way that I think about this area is... My current belief is that there isn't a strong set of things that's really backed by research that I'd be excited to go rolling out now, but there is room for more research to find more such things. And so, for me, it's in the research bucket. And I haven't talked about research yet, but like I said, scientific research broadly, which could include social sciences, is like one of the four buckets we think about. And that's actually, within social sciences, that's one of the topics that I find most interesting and want to start looking at when we have the capacity is like, is there room to do more and better work on that? But, it's not something we've gotten to yet.

27:44 HK: And just a quick note on timelines. Our focus this year was policy and global catastrophic risks. So, we've got pretty developed views on the best causes in those. And what we're trying to do right now is, I think we're at the point we've largely looked at all the causes that are most interesting to us, and we have to decide which ones we're going to work on and how much to

work on them. And that's going to dictate kind of staffing up the Open Philanthropy Project over time.

28:11 HK: And then the causes, the areas we haven't looked as much at are scientific research and the other one, global health and development. And those are just more, we didn't have the time. We didn't have the staff. We put some thought in to them this year, and especially on science, put a decent amount of time in and definitely have things to say, but I think they're further away from the point where we can say, "Yeah, we've looked in to them. And we got to think causes now." So, I think I'll just go ahead and do a little spiel on global catastrophic risks. Then, I'll take some more questions for a bit.

28:41 HK: So, I mentioned, we're interested in things that could kind of threaten global civilization. This work has been led by Howie Lempel who's not here today because he's actually going to a bunch of meetings on pandemics and research that can pose risks and things like that. So, a quick rundown of what... Again, I'm going to highlight the things that we're most interested in and also talk a little bit about the criteria of how we think about things.

29:08 HK: Again, it's going to be a familiar framework. We look for issues and problems and dangers that are... I guess, in my head, I usually use the word scary instead of important, but it's the same basic idea. Things that you could just imagine having a decently high probability of doing a really terrifying amount of damage sometime in the near future, so scary. Tractable meaning there's something you could imagine doing today to prepare for it. And then the classic, uncrowded. And so, in terms of some of the stuff we've looked at, it's probably easiest to start by talking about climate change. Climate change is probably the best known global catastrophic risk. It looks like we're on pace to raise the planet's temperature a good amount or a bad amount in the next several decades. That would have a lot of really bad effects. It would have really major costs. This is an issue that gets a ton of attention. It's actually one of the only global catastrophic risks that I think gets a ton of attention from philanthropy and non-profits.

30:06 HK: Something that I think is interesting though, is that if you look at the existing projections of what's going to happen, the best guess of what's going to happen with climate change; it's really, really bad, and it's especially really bad for the global poor, but it's not so bad that it leaves us worse off in the future. Actually, the projection is that the world as a whole and on a per capita basis is much wealthier a 100 years from now and much better off in most ways; even then of climate change. The thing that I'm most scared of with climate change is if those projections are wrong. And it's hard to have much confidence in them either way. And so, things could turn out to be much worse than we think.

30:41 HK: And so, a lot of what I think about with climate change is what happens in the worst case, and what could we do about that. And when thinking about that, I think the issue of geo-engineering becomes naturally salient. This is... Geo-engineering is the idea of purposefully interfering or manipulating the planet's environment to offset some of the harmful effects of climate change. And that could include things like shooting sulfate aerosols into the atmosphere to change the way that sunlight is reflected and actually cool the planet. This is something I would put in the category of like sort of a crazy dangerous thing to do with potentially really bad consequences that we may want to do anyway if the alternatives are bad enough. And that makes it something that to me might be a good idea to be putting a lot of resources into understanding as much as we can, because I think the day may come...

31:29 HK: The other scary thing about geo-engineering and especially sulfate aerosols and other forms of sort of cooling the planet by changing sunlight reflection is they're actually pretty cheap and fast, which makes them very good for a worst case scenario. It also means that any really... I mean physically speaking, anyone in the world can decide they would like to cool the planet, and then they can cool the whole planet. That can cause a big problem. So, I mean, two of the things we think about is one, research. How can we understand the safest most effective ways to do this, and what would the downsides be, and what would the risk be? And be governance; who's allowed to do these experiments, who's allowed to do these interventions, what norms can we agree on internationally, what arrangements can we come to that reduce the risks of a giant conflict breaking out over unilateral geo-engineering?

32:15 HK: So, this aspect of climate change, I think, gets very little attention, and if you look at who's funding it, you look at who's working on it, it's actually a pretty thin space. And it's very thin compared to the whole world of people trying to put a tax on carbon and trying a lot of different angles to make that happen. So, this is a good example of where... We look at climate change as a whole or a lot of sub-aspects of climate change, and we say, "This is really important, but also it's very crowded." But in some ways the most crucial part of the space that we've seen, the part that's going to be most relevant in the worst case, is actually pretty uncrowded. And so, that becomes a cause that is interesting to us, and it's something that we're actively working on, and may make a grant in sometime soon.

32:57 HK: Other causes we're interested in, so buyer security is a big one. The Ebola pandemic has been really bad; could still become much worse. But actually, I think, there are certain ideas of flu pandemics that could be much, much worse than anything Ebola's realistically going to turn into. That's kind of any-day kind of threat, and the better our everyday defenses are, the less of a risk it becomes. There's also a less everyday, more kind of sci-fi future speculative threat, which is the idea that at some point, we'll be able to manufacture and synthesize pandemics that may be worse than anything that arises in nature. That is probably one of the leading candidates. If something's going to wipe out the entire human race in the next 100 years, I think most people would say that's one of the leading candidates to do it. And in some ways, the two overlap a lot in terms of what you can actually do about them.

33:49 HK: So, for example, strengthening routine surveillance is in some ways a protection against both kinds of problems. And one of the interesting things we've found is we've looked at the Ebola situation is that people are now getting more interested in strengthening routine surveillance, and so there may be opportunities to do something. So, we think it's an important issue. We think it's tractable. In terms of crowdedness, there's a lot of government resources that are working on things related to pandemic preventions, surveillance, bio-security, but there's very little philanthropy. And so, one of the big questions we've asked is "What does that mean? Does that mean there's good opportunities or not?" By default, I would generally think philanthropy in government have different strengths and weaknesses.

34:31 HK: And so, if you see one and not the other, you see government without philanthropy, there may be opportunities for philanthropy to improve the way government is spending its money. And that could be a really good opportunity. Preliminarily, that is kind of what things look like, and we think there will be opportunities to do things, and we're looking into it further. Some other causes we're looking at a little bit more briefly, artificial intelligence; this one is, again, more

speculative. I think that the odds look decently good that we could have extremely powerful artificial intelligence in the next few decades. It's definitely not a certainty, but it's a possibility. And if that happens, that can create a whole bunch of new problems that we've never had before. I mean, a problematic coding error could definitely take on a whole new meaning at that point.

35:19 HK: And our big question here is, what is there to do today to actually... Is there anything useful we could do today to deal with that kind of risk and make that less dangerous in the future? And I think the jury is still out on that one, and I expect to know more in a couple of months. And then... I mean, because of the work we're doing, not because of... [laughter] And the work we're doing is not building AI, so... Then another cause is geomagnetic storms; these are like space weathers, solar flares, that are believed to pose a threat to the electrical grid. We saw a lot of disputes over just how dangerous they are, and how much of a blackout they could cause and over what area. It's one of these things where if you decided it was a big risk, it would be an imminent risk.

36:01 HK: It would be a pretty large probability risk 'cause these storms happen. And there have been some really big ones before the age when it would have been really bad. And there are things we could do about it. So we're trying to figure out how big the problem really is. And that's been another project David Roodman has been working on. And he's kind of the person we go to a lot of the time when we need to sort out what a complex and difficult literature says about something. So we expect to have that report out. We're also interested in nuclear security, which I would say is the other risk that gets real attention from philanthropy, but is a real risk and an imminent one. And we also think about general things we could do to make the world more resilient, more robust, to catastrophic risks. So things like food security. A lot of these risks, the biggest threat they pose is what they would do to the food supply and the agricultural system. So what could you do to protect against that? So I'm happy to take questions on those or more questions on policy or questions on Open Philanthropy Project as a whole. Yeah?

37:01 Speaker 7: How much do you find interactions between different [37:07] _____ causes? Possibly in different areas. So, the example I can think of is there was a paper on how more migration would significantly mitigate the negative consequences of climate change to the poor?

37:19 HK: Yeah.

37:19 S7: Two of your issues that are completely unrelated, seemingly, except that they really significantly might actually interact. Maybe there's more of those.

37:27 HK: Sure, so how much interaction do we find between the causes? For example, more migration might mitigate the negative effects of climate change. We see some interactions. It hasn't been a huge part of what's been going on. I think usually when we do see an interaction, it's often a speculative one, and it usually goes in the right direction. In other words, it doesn't make us want to work on a cause less. The example you gave is an example of that. But, yeah, we keep an eye out for that stuff. And I think, one of the things that I hope will eventually happen is that by being so broad and looking at so many different places, we see opportunities that others can't see because of combining things. And so one of the things that I'm currently pretty interested in is we've been looking at scientific research. We've been trying to understand how policy change works. The intersection of science and policy, like science policy. I think is a space that...

[laughter]

38:22 HK: Thank you. That's a space that I think is pretty interesting, and I think in some ways, I feel like the best science funders are not really into policy. And the best policy funders are not really into science. And so if you were into both, there might be really interesting things to do. Yep?

38:37 Speaker 8: How do you decide how many people to staff on Open Philanthropy?

38:43 HK: How do we decide how many people to staff on Open Philanthropy? That is a multi-dimensional question. And that kind of gets to our general approach to hiring and our philosophy of hiring as well as our ability to plan. So generally the question comes down to, what is the work that we need done, how much do we know about the work that we need done, how able are we to predict who's going to be able to do this kind of work, and how able are we to hold people accountable and evaluate how well they're doing this kind of work? And that's all a prerequisite to asking who can do this kind of work? And so, I guess that was a confusing answer, but I guess it's... I think there are early in Open Philanthropy Project, it was basically just me and Alexander. We were the only ones working on it, and it was feeling our way along and not having much of a plan. And especially for some of the causes that we knew less about, like for science.

39:39 HK: That was just... I'm usually the one who works on something when we have no idea what we're going to do, and we need to come up with a framework or a basic approach. And it's as we get better defined, and as we say, "Okay, we know that we want this work done. We know that we need to look into this cause by speaking to five experts, or several experts and get a picture of how important, how tractable, how underfunded it is," that's a better defined problem than what should we do in science? And so then we're able to think about who's able to do that. And then we look at who we know through our network, who we can find, who is already on staff. And whenever we can make a match between a person and a job we need done, we try to make that match. And so it's somewhat chaotic for that reason.

40:27 S8: I guess I was more thinking, you're spending your time. You decided to do Open Philanthropy at some point, instead of spending that time on GiveWell, researching more charities. I don't know. How do you make that judgment call?

40:37 HK: Yeah, how do we make the judgment call of spending time on Open Philanthropy Project versus top charities? I think that has largely come down to just personal preferences or views about how to do the most good. And I think when we started Open Philanthropy Project, which was then called GiveWell labs, I think everyone on staff was in agreement that this was something we should do. That GiveWell as an organization should do, that someone should work on it. That it was a huge opportunity. I think everyone was also in agreement that we had to keep GiveWell top charities going. That that was a really good project and a really important project. And then we sat down and talked about who's capable of working on Open Philanthropy Project, given how early it is and how little guidance we have to offer? Who's capable of leading the work on GiveWell top charities and making sure that still stays alive and stays good and gets better? And how can we allocate our staff so that we get the best possible results on both? But I also think, ultimately, we envisioned them becoming separate organizations. And yeah, I think that in some ways, personal preferences and personal values come into the decisions as well.

41:47 Elie Hassenfeld: I think the place where Open Philanthropy is today is where GiveWell was a

long time ago. And the goal with GiveWell is to hire people and bring them along so they can do that without Holden, without me, without other people that today are in senior roles, and I think that is the... That's seems like the intuitive way to try and run an organization to try and make it go beyond just the people who are literally there today. And right now, Open Philanthropy is still so early that I think there's a much smaller group of people who can meaningfully contribute to it.

42:21 S9: I got you.

42:25 HK: Yep.

42:25 Speaker 11: In terms of global catastrophic risk, what's the interaction between GiveWell and Max Tegmark's group on the east coast as well as the groups at Cambridge and Oxford?

42:34 HK: Sure, on global catastrophic risk, what's our relationship to Future of Life Institute, I think it's called. That's Max Tegmark's group in the US, Future of Humanity Institute in Oxford. You didn't mention Machine Intelligence Research Institute in the Bay but... Yeah, the usual answer to this kind of question is going to be like we talk to them. I think in this case, we're quite friendly with all those groups, especially because with global catastrophic risk and with those groups in particular, those people are really into effective altruism which is the same thing we're into. It's this general idea of doing the most good you can with every dollar and every hour that you have. According to a set of values that is sort of saying "All lives across the globe count. We want to help them all as much as we can." We're not, it's not really driven by equality is a fundamental good; liberty is a fundamental good. It's driven more by well-being. So this is how I think about the effective altruism platform and the effective altruism value set, and those people are really into that idea, and so we tend to be really friendly with each other.

[pause]

43:41 HK: Yep.

43:43 Speaker 12: Could you say more about why policy makers over value inflation relative to unemployment?

43:50 HK: Maybe Alexander could say... First, I'll repeat the question. Can I say more... Can we say more about why policy makers over value inflation relative to unemployment?

43:58 Alexander Berger: Yeah. I think it depends on which policy makers you talk about. So the particular institution that's most responsible for making monetary policies, is the Federal Reserve which is the US' central bank. If you talk to people who have been involved in the Federal Reserve in the past, especially people who weren't already economists in 1970s, so anybody under 60. What you'll hear is that the Federal Reserve live and dies by the reputation of Paul Volcker who is widely perceived as being responsible for having crushed inflation in the late '70s, well, really in the early '80s. And so the Federal Reserve... According to the people we've talked to, views that as their greatest victory, the victory over inflation. And so, today, when you have, not a problem of inflation, but actually a problem of unusually high unemployment, although that's today, literally today, not as true, it has been really true for the past five years. There hasn't been the same emphasis on fighting unemployment that there has been on continuing to have the victory over inflation that you've had for the last 30 years. And so I think that the standard is what fight do

people remember and what do they see as their core institutional goal, and the Federal Reserve, I think, sees it as inflation partially because of this historical factor.

45:19 HK: You have a follow up on that?

45:21 Speaker 14: A related question.

45:22 HK: Yeah, sure.

45:23 SP 14: About macroeconomic stabilization, there have been people that are talking more about guaranteeing their own income as a way of dealing with the possible large threat for unemployment as automation becomes more and more of an issue? Is that something you considered in your research?

45:38 HK: Yeah, have we thought about a guaranteed minimum income as a response to the threat caused by unemployment as technological automation, right, becomes more of an issue? Yeah, that is actually on our list. It's a little bit lower down than the causes I mentioned, but it is a cause of interest. And I think, I think there's a couple of things to say about this. One: There have been false alarms before about the coming technological unemployment. It's certainly not something we know is coming. It's something that might be coming, and I think it's worth thinking about what good policy would look like in that event. I don't think it's obvious that the right answer is a universal basic income. In other words, what happens, let's say, in a hypothetical world, possibly a future world, where a very large percentage of the population is essentially unemployable because there are automated ways to do any economic value adding job they can do.

46:29 HK: Should you respond by everyone having a guaranteed income and not having to work... And there's people, I think, on both the left and right who are really worried about that idea, who say that it's good for people to work, and that this kind of set up is inferior to some kind of set up that encourages work or subsidizes work. Another way you could have this, you could subsidize work heavily such that people are basically doing jobs that might pay incredibly low wages, but they take home perfectly good wages. And that may be a different kind of structure so... I think that's actually a debate that could be a higher quality debate than it is and making debates more prominent and higher quality than they are is one of the things that I see as like in the philanthropist tool box, or one of the things a philanthropist can try to do. So I think it's an interesting cause to work on. I do think it's a pretty speculative and may never come to pass, and I'm not sure that it's as strong as some of the other causes I've mentioned, but it's definitely an interest of ours. All the way in the back, yeah.

47:28 Speaker 15: The GiveWell approach that you all describe seems very clearly unique, I'm wondering given that you're in the early days of your Open Philanthropy Project, are you finding that your approach is significantly different than, say, a big foundation who decided they wanted to attack criminal justice reform and would look at the four buckets of policy and organizing and...

47:49 HK: Right.

47:50 HK: What are you finding so far?

47:52 HK: Yeah. Are we finding that our approach on Open Philanthropy seems unique? Does it

seem different from other foundations? Does it seem the same? We've definitely seen similarities, and I think to some degree, we're learning from them. To some degree, we're going to them and saying, "How do you attack a new cause? And how do you figure out what to do and figure out whom to grant to?" And we have learned a lot from other foundations. I think there are three ways in which we're pretty unusual, that I think are fairly robust, and that these are the same three ways I guessed we would be unusual when we first started the project. So one of 'em is... I'll kind of go in decreasing order of how strongly I feel they're unusual.

48:30 HK: One of them is the transparency. So this is the one that's just definitely a value added of ours, definitely something different, where you go out and you say... As Carey did basically, "I am trying to do a huge amount of good by giving away a good amount of money. What do I do?" And it's just hard to figure out what are even the debates, what are the questions, what is the information, what are the arguments about how to do this well and how to do it poorly? There's some stuff out there, there's very, very little. And I think if more groups were kind of just making an attempt to document their thinking. We can't do it as exhaustively as we do with GiveWell top charities, but we try hard to explain why we're doing what we're doing and to have it all there so that someone who wants to learn from us can learn from us. So I think that's the number one way in which we're unique.

49:18 HK: Number two, which I think is also pretty strong and pretty robust, is the effective altruist attitude of, "I want to work on whatever does the most good." So we've just had a million conversations where we go to a funder for advice and we say, "What would you do in our shoes?" And they say, "Well, first you have to decide what cause you're passionate about." And this is what philanthropic advisors tend to do too is they'll say, "We work with you to figure out whether you want to work on education in St. Louis or homelessness in San Francisco and then... And we go from there. And our kind of consistent response has been, "Well, we're excited about working on whatever cause is best to work on and will do the most good".

49:54 HK: And I think that attitude and also the particular value set of effective altruism, of just how we weigh well being and global issues, and the weight we place on economic development as well, I think those are unusual. And so, in some ways, it's been kind of hard to find kindred spirits on that, and that's been a source of frustration especially on policy where it's very good to find people who share the broad platform you have and agree with you on all policy issues. That's been hard.

50:24 HK: And then, I think the third way is just, we have this particular way of breaking down problems, of analyzing why we think what we think, of poking it, of using empirical evidence, and it's kind of an epistemology that comes from the top charities work that we're adapting to the open philanthropy work. So we've dropped any requirement that things be supported by randomized control trials but we're still constantly asking, "How could we be wrong about this and what can we do to find out whether we're actually wrong about it?" And I think that's something that is... Certainly, the particular way we do it I think is somewhat unique, and the extent to which we do it I think is somewhat unique.

51:04 S?: Do you want to say a little bit about different ways to be in a cause area and how we're thinking about that?

51:09 HK: Yeah. One of the big discussions we've been having a lot recently, which I think is an

example of something that it's hard to even know this is a question if you're just getting into philanthropy, and it's certainly hard to find much intelligent debate about it, is what should you do? How many causes should you work on and at what level, and structurally how should you do it? So there's these different schools of thought. I think a very common thing for a foundation to do is they hire program officers or program directors, and those people will work on one cause, and they'll spend all their time on it, and they're often experts on the cause before they come in. So you're working on criminal justice reform, you hire someone who's extremely experienced in criminal justice related issues, and they spend all their time in criminal justice reform for years. This has a lot of advantages because there's a lot of nuance in that space, and there's a lot of relationships and there's a lot to know. And so, having someone who really knows their stuff, I think, in some ways, is the only way to really understand what grants you're making.

52:10 HK: There's disadvantages too, and I think one disadvantage is potentially inefficiency. It's one person per cause. Another one is that you have to make a long commitment to a cause, and you have a person who's, their job is to get as much money as they can spent on that cause. And so, you kind of have to give them a budget. We've been looking at funders who use other approaches who maybe have a lot of things they're interested in but are very opportunistic and look for opportunities in ten different spaces or ten different ideas. And then, when they see one that's really exciting because the organization's exciting or the person they're funding is exciting, they'll then kind of focus in and spend a very long time understanding that idea or that person. So that's like a pretty different approach.

52:54 HK: I can name what I think are some interesting cases of philanthropists who seem to have done things that are about as impressive or more impressive than the big, traditional philanthropists, and they did them with having much less staff and much less expertise. And I think it's kind of interesting and that's something we'll be writing about more on the blog. But that's... Asking how that happened is one of the questions that we're A, tackling with the history of philanthropy project which is where we're trying to create better information on what philanthropy's done well in the past. And B, it's kind of inspired us to ask, are there better ways to do philanthropy or at least different ways that have different strengths or weaknesses? Do you really need the level of expertise that comes with having one person on a cause or could we go further by taking ten causes of interest and looking for only the very most outstanding opportunities we see? So I think there's big pros and cons there and big debates to be had. And part of what we're trying to do is get the right answer and part of what we're trying to do is just make the question bigger and more prominent.

53:55 HK: We only have time for like one more question.

53:56 HK: Yeah. Okay. Now, I have to choose. Someone who hasn't spoken as much, yeah, you.

54:04 Speaker 16: Do you think at all about sort of a meta-cause of trying to make more people aware of and interested in effective altruism?

54:10 HK: Sure. Do we think about the meta-cause of trying to make more people interested in effective altruism? We certainly think about it. I think our attitude at this point, there are other groups trying to do that. I also think in some ways, those groups are decently positioned to get funded. There's a set of effective altruists, a lot of effective altruists gave away a lot of money. And they are able to recognize effective altruists. Like in some ways, the challenge for us is helping effective altruists donors find things that would have never been on their radar like organizations

working on pandemics who don't hang out with effective altruists and talk about effective altruism.

54:48 HK: So, in some ways, I feel like that's a more productive place for us to be. And I also just... I think in a lot of ways, we feel like our role in the effective altruism community is to make things more concrete, and tangible, and real. So, it's one thing to go around saying, "You should do as much good as you can." But then people will say, "What does that mean? What do I do? Where do I give?" And that is, I think, the question that GiveWell really tries to specialize in answering. So, I don't think we just want to turn this into an infinite loop. I think there are groups working on this. I'm really glad there are. I think the best thing we can do for this set of values is to find concrete, tangible ways to accomplish things by those values. And I think that will also help the community grow.

55:35 HK: So, we're going to wrap up here. Thanks to all of you for coming out and engaging with us on these questions. I know it's not the standard charity event that you might go to. But it really helps us to have people who want to come and grill us and ask us questions, so we really appreciate it. I think a bunch of GiveWell staff will probably stick around for a little while if people have more questions they want to ask one-on-one, but thank you.

55:59 HK: Thanks.

[applause]