

GiveWell New York Research Event, December 4, 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.

0:00:00 Holden Karnofsky: And I think we mostly have the same people here who were here at the first session, so I won't give a lot of intro overview. This is the second session. This is going to be me talking about the Open Philanthropy Project. And again, I just want to emphasize that I'm happy to hear any questions no matter where you're coming from or starting from, because we definitely don't expect people to read all the stuff we put out. And we offer these events largely as an alternative to reading so much stuff, and certainly not as the next class or whatever. So, I'm going to talk about the Open Philanthropy Project. My basic plan... First, I'm just going to give the overall view of what the Open Philanthropy Project is, how we're doing things, and then I'm going to talk about our work on US policy, and that's going to be the longest part of this session. And then, I'm going to take questions for a little bit. And then I'm going to talk about our work on global catastrophic risks, and take some questions. Then, I'm going to talk about our work on science and other causes, and take some questions.

0:00:59 HK: So, basically, GiveWell started in 2007. At that time, we were trying to serve first people who were in the situation of saying, "I have some money to give, I don't have a lot of time to think about it. I want it to go to something I can understand." And that's how we came to the idea of promoting sort of proven, cost-effective charities. In 2011, things changed when we got a lot of interest from Good Ventures, which is a new foundation founded by Dustin Moskovitz, the co-founder of Facebook, and the president is wife, Cari Tuna. And they were basically just getting started in philanthropy and saying, "We're going to be giving away a lot of money over the course of our lives. We're going to be at some point, probably one of the major foundations in the US, and we want to do as well as we can. And we're young for people who are thinking about philanthropy. We want to do it as well as we can."

0:01:53 HK: We met them, and felt like we were really interested in a lot of the same questions, and decided to work closely together, and jointly kind of launched this project that at the time was called "GiveWell Labs". Now, it's called "Open Philanthropy Project". And the idea is to ask the question, "If I'm a major philanthropist, and I can hire my own staff and make major grants, and create organizations, and take big risks, and make long-term bets, and basically do whatever, how do I do as much good as possible with my giving?" So, it's largely continuous with the mission of GiveWell. It's asking, "How do I give GiveWell as possible?" And then, it's trying to be very transparent and open with the world about how we're thinking, so that people can learn from it, and so that we can influence the way others give.

0:02:39 HK: The difference is that the criteria are much more open-ended, and the target audience is a little bit different. So, we're mostly designing this as being targeted toward major donors like people who could be giving away \$100 million or a billion dollars over the course of their lives. We are, I think, at the process also finding opportunities for individual donors to contribute to, but that's not like the thing that drives the project. So, that's Open Philanthropy Project. We've been... Basically, this project is much newer and more experimental than GiveWell. It's important to recognize that. When GiveWell started in 2007, it was me and Elie. It was two people, and we were experimenting a lot, and it took us a long time to get to where we had, good top charities.

0:03:27 HK: That's kind of where Open Philanthropy Project is now. It's been two years of pretty intensive work, but it's been a small part of our staff working on it. It was basically me and Alexander Berger for a while, and we've recently added a couple of more people, plus Cari the President of Good Ventures. And so, it's an early project. We don't have been giving recommendations yet. And because of the size of the question, and the ambition of it, and the open-endedness of it, and how much capacity we're going to have to build, it could be years before we're saying, "Give here." We see this as a really long-term project to build a lot of capacity to move a lot of money, to really big long-term and potentially risky stuff, and that makes it a very different project from GiveWell.

0:04:13 HK: Basic intellectual framework we're using is also different from GiveWell. So, GiveWell top charities looks at the level of the organization or the intervention. We talk about bednets, we talk about iodine, we talk about the Against Malaria Foundation. The main thing that we've talked about for Open Philanthropy Project so far instead is what causes to focus on. And by a cause, I mean, basically a problem in the world that a philanthropist might be able to do something about. And we've written about why we've kind of come to this lens, and why we think it's the right framework. But basically, it's hard to evaluate an organization or a project without a lot of background knowledge about what they're trying to do and who else is trying to do it, and how well the pieces fit together, and what most needs to be done, especially when you really can't rely on randomized control trials the way we do with GiveWell top charities.

0:05:09 HK: A lot of the information about how to give well is going to be based on knowledge about people, and about organizations, and about political dynamics, and about scientific issues. And so, when you need so much background knowledge to assess one project, and then that background knowledge helps you assess a whole bunch of related projects, it seems like the most efficient question asked is, what background knowledge do we want to develop? What do we want to be really intense about? What do we want to be knowledgeable about? What do we want to be experts in? Another reason I think this way is that a lot of what major philanthropists can do is instead of just finding existing charities, they can create new opportunities. They can... A lot of times what we found is if you go around saying, "I want to fund X," new opportunities will come to you that you could have never found by just kind of asking people what to fund in an open-ended way.

0:05:57 HK: And we've got some interesting examples of that. So, that's what we're trying to do right now. We've kind of divided the universe of different things you could do into four very broad categories, and then within those, we divided it further. So we're... One of the broad categories of philanthropy that we look at is policy advocacy in the US. And what I mean by that is basically, trying to influence... It's actually... A lot of not is advocacy, but it's trying to influence the direction that policy takes in the US by putting out information, by helping people understand the impact of different policies, by supporting groups that may put pressure on people to pay attention to certain interests. This is... A lot of the major foundations today will tell you, this is how to get the most bang for your buck. Because by spending a small amount of money on... People who try to put an issue on the map and try to raise the profile of an issue, you could change a lot. You could help get a lot changed or help increase the probability that a lot changes. And government's power and money is so much greater than what any kind of foundation has, that, that becomes a very leveraged impact.

0:07:08 HK: Another general area we're interested in is scientific research. So, one of the things

we've done is look for all of the examples we can find of really amazing, successful, philanthropic success stories. And there's not a lot of work been done one that, but what we've seen, a lot of it is... Again, you can get a big return when you spend a little bit of money funding some sort of investigation or research that then leads to a new insight that the whole rest of the world can use forever, for free. So that's another way to get leverage.

0:07:39 HK: A third thing that we think about for having maximal impact is global catastrophic risks. These include things like climate change. They also include lower profile issues like pandemics or solar storms that may wipe out the electrical grid, basically anything that kind of threatens the whole world at once with a very low probability. It's not really clear whose problem these things are, and it seems like a good fit for philanthropy, because you can have a very small impact of preventing something really, really large and catastrophic that could derail or alter the course of civilization. And I think in the case of something like climate change that's not hyperbole, and even in the case with some of these lower profile risks, that I think it's not high hyperbole either. So that's a chance for big impact, is to raise the profile and raise our preparedness for an understanding of global catastrophic risks.

0:08:33 HK: And finally, we remain interested in global health and development as a focus area, because of the thing I said before about your money going furthest in the poorest parts of the world. And we don't just want to do that in the normal top charities way. We want to think broadly about what are all the things a philanthropist could do to help the developing world.

0:08:53 HK: So with those out of the way, the basic way our agenda looks is, this year we've been very focused on the first two categories I said: Policy and Global Catastrophic Risks. And we've had a goal of deciding what our focus areas are going to be, like what we're going to build knowledge about, what we're going to build staff for. By the end of this calendar year, which is obviously just about upon us, and I think we're going to hit something like that goal, 'cause I think we largely have done the research we need to do, and largely do know what we need to know to think about that decision. And then the other two categories are more for the future. So it's like, once we pick causes in these areas and we have the capacity, we're then going to be putting more effort into exploring scientific research and global health and development, and picking focus areas in there.

0:09:42 HK: So that's the overall framework. Next thing I'm going to talk about, and then I'll take questions, is I'm just get into a little more detail on work on US Policy. And I think that'll be helpful in just thinking about how this whole thing works, because I think, obviously everything I've said so far is a lot vaguer and more abstract than the GiveWell top charity stuff. So, when we look at US Policy, the people working on that right now, first off, there's Alexander Berger, who's been with us for a while. He's a Senior Research Analyst, and then Shane Astrom, who's a more recent hire, is actually in DC, and is someone with much more of a political background. And this is an example of deciding on the things you want to be really good at, and part of how you handle that is you hire people who can be really good at those things.

0:10:26 HK: With policy, we've tried to canvass the universe of US Policy areas, and we've looked... We've asked three questions about most of the areas. We've asked, number one, how important is this policy? So, there's... People would like healthcare policy to be better. People would like gun control policy to be better. How important is each of these? How many people would we help by changing policy the way we'd like to change it?

0:10:53 HK: Question two is, how politically feasible is it to actually make a difference here? And so I think there's some issues that are extremely partisan and divided, and it's been a decade of people fighting each other and nothing seems to be happening. And there's other issues where there may be some sort of window of opportunity. I think it's easy to overrate this criterion, and most of the actual experts we've spoken to on policy have said, "Don't get too excited about what looks politically doable, because what's politically doable changes drastically and unpredictably on the relevant time frames." And so, a lot of times, the worse thing you can do is chase headlines, get into some area because it's the news, fund a bunch of groups, and by the time your funding is making any difference, the whole situation has changed. So, that's something we've taken into consideration.

0:11:43 HK: Finally, classic GiveWell question, how crowded is the area? So, if we're looking at climate change policy, well there's a lot of people working on climate change policy, and a lot of money going into trying to improve our policy to reduce carbon emissions. There may be other causes that are comparably important, and don't get as much attention. And what we're really looking for is a policy area that is not getting enough attention for how big a deal it is. That's the summary of what we're going for. So now, I'm going to list some causes that we think are particularly outstanding by one or more of these criteria. And usually, the way this works, I mean we don't have any precision here, it's not like we can say, "This cause is a 2.2 on tractability and a 3.3 on importance and we're going to multiply them together." More what we do is we kind of look for the most important causes, and ask which of those do best on the other criteria, and we look for the most ignored causes and ask which of those do best on the other criteria, etc. So one cause we're interested in, as an example, is criminal justice reform.

0:12:45 HK: This is the one we've done the most work on. We're currently working on a deep dive of really trying to get to know everyone in the space, all the issues, think about all the things we could do and potentially even hire a specialist. Criminal justice reform, so the US has incredibly high levels of incarceration, and incarceration is really inhumane and really messes with people's life prospects, and there are good reasons to work on that. It's not the biggest issue in the world in our opinion. There are policy issues that matter more. And it's not a neglected issue. There are funders working on it. What makes this area appealing to us is that it is one of the few areas and kind of the only area, though there's a couple other candidates that we've seen people point to again and again, and say, "There's opportunity here."

0:13:35 HK: There is political opportunity here because, and long term political opportunity, not just a one day news headline, because there's a structural situation where crime is lower than it's been for a very long time. Incarceration is very high and state budgets are really tight. And so where 10 years ago, everyone, or definitely 20 years ago, everyone was competing to be tough on crime and the only people talking about less incarceration were on the left. Now you have people on the right and the left, actually collaborating. And some unusual people are kind of joining the cause. So you know there's right on crime which often, if you go there, you'll see like Grover Norquist on the front page and they're saying, "We need to stop locking so many people up."

0:14:18 HK: You can argue about why they're saying that. One reason I think is that it's a lot of government and it's a lot of money. At the same time, it's also really inhumane. So there's a political opportunity here and that makes it an interesting cause. It's also a really something we kind of realized about it later, it's what we call a "high venue space" which means there's a lot of different

places you could go to try to make change. So if you're working on, there's a lot of policies where the only thing you can do is hang out in DC and hope to influence federal policy. Criminal justice reform is mostly a state and local issue. You can influence judges, prosecutors, states, municipalities, there's a lot of opportunities to be opportunistic and find opportunities to make a difference.

0:15:03 HK: So it's a big space. It's a complicated space. There's a lot of funders in it. But we're trying to explore it very thoroughly and we have a whole bunch of ideas about where the biggest juice might be. Another cause, which is very different that we're interested in is labor mobility. So the argument here is that the best anti-poverty intervention that we are aware of is people moving from a poor country to a rich country. Actually stepping across that border can cause someone's income to go up by a multiple. And it's just stepping across the border. It's the same person in a completely different economy with a completely different price on their skills. Obviously, this is a very politically-heated issue. This is kind of the opposite of criminal justice reform.

0:15:46 HK: It doesn't look like it's going anywhere anytime soon. What we think is interesting about this cause is that even a small win in political terms could be really, really huge in humanitarian terms. If you look at how much you're causing more money to go to people who are coming from extremely low incomes. And so even small wins can be very big. And also most of the philanthropic attention in this space is focused on the people who are already here, their path to citizenship, and not on this fact that crossing the border in many ways is the most important part of this whole story. So that's an area we're interested in. Center for Global Development is a big proponent of this idea and we've supported some of their work on it.

0:16:29 HK: And by we, I mean, GiveWell and Good Ventures working together as part of the Open Philanthropy Partnership. And we've also supported a couple of other groups that we found are really, there was no other support for them because they're so focused on this labor mobility issue as opposed to immigrants' rights issue. Then there's macroeconomics stabilization policy, which we're trying to keep up with GiveWell top charities in having unsexy, complicated names for the things we do. Macroeconomics. So basic story here is, there is a lot of debate over what happens when there is a recession and why recessions happen, and what the government can do to make recessions less severe. This is another area where I think there's a lot of reasons to believe this is a tough space to work in.

0:17:19 HK: It's also incredibly important. If you could make, slightly decrease the depth of a recession in terms of how that compares to other causes where you might get a win, the win would be really huge. And also there are a ton of economists and intellectuals who work on macro economics, but in other ways the field is very under-appreciated and kind of neglected. There's very little in the way of groups doing public advocacy on this issue, trying to be political about it, even trying to come up with clever new ways to design policies that would be both good for the economy, but also like actually passable. And we think there is some fertile ground for that, to try and think about these things 'cause it's not really a strict left and right issue.

0:18:02 HK: So that's another area we're looking into. We've made a few grants on it. We could do a whole meeting on any one of those grants, so I think we're going to have to be selective here in what we talk about, but that's one we think is interesting. Just to rapid fire through some other issues, we've thought about a little less, policy around land use zoning, how land is used. Not the most important issue in the world, but there's practically no one working on it that we can find. And

so that's like an interesting question, is what happens when you see an issue that no one is working on? Can you get people to work on it? Can you raise its profile? Can you get some state and local wins that for the amount you spend, there might be actually a really huge return on your money. Foreign aid policy, we're interested in. I actually think there's a pretty strong infrastructure already around advocating for more and better foreign aid, but we're interested in finding more stuff we can do in that. Factory farming, very popular topic with our audience, very controversial. I think animals are treated really, really badly on factory farms. I think most people don't know this or don't know the extent of it, and so the question of whether that could change also like a little bit less partisan in a lot of issues, arguably very important depending on how we exactly value animals and now particularly crowded.

0:19:20 HK: There's a ton of animal charity, but it doesn't focus on this issue. It focuses more on strays and things like that, endangered species. What are we actually going to do? What are we going to focus on? There's this whole other question that we've been grappling with, which is how many causes should we work on? How should we work on them? How many people should we have on our team? What should those people do? And there's obvious debates that I found very interesting, but they get very little attention about just how to structure philanthropy. So I think the most traditional way that a big foundation operates is they'll pick a number of issues.

0:19:54 HK: So they'll say criminal justice and labor mobility, and they'll hire more than one or one person per issue and that person will spend all of their time just thinking about that issue. And they'll get to know everyone in the issue, they'll often will be a pre-existing expert on it, like a criminal justice expert and that's what they'll do and then they have their program areas. But I've also seen some other groups that take a very different approach. They might be interested in 20 things at once, and they aren't experts on any of them, but they know a little bit about all of them, and they're looking for an outstanding opportunity. They're looking for something that really pops or an organization that's really got great leadership and needs money and that's a hard combination to find so they grab it where they see it.

0:20:35 HK: And which of those models is more promising and is going to make more of the money? I think we have yet to learn what's going to work best for us, and our skills, and also just works best in general. Where we're leaning right now is trying both of the things I just said. So I think we want to take at least one cause and it's pretty likely to be criminal justice reform where we take the... Try to find an expert or a real specialist approach, and then we may have more of a GiveWell generalist, someone like Alexander work on a bunch of other causes at once, and look for the best opportunities he can possibly find. Other challenges here, hiring is always difficult and now we're talking about how to build a team, not just how to pick an organization. So a lot of what actually happens in the next few years, it's going to be driven by, we're going to be trying to find the best people and the best organizations in a lot of causes, and a lot of what we actually do is going to be driven by whom we meet and how promising they are and what we end up jumping on. So that's where we stand on politics at the moment. I'll pause there and take questions before covering some other topics. Yep.

0:21:45 Speaker 2: Thank you. Bob Kuttner. I'm interested that you put macroeconomic stabilization on that list, and the reason is, that's a policy area that where so-called expert opinion is ideologically polarized. You have a lot of center-right philanthropy, Peter G. Peterson Foundation, to give you exhibit A, pouring money into macroeconomic stabilization, defined as that you're stealing its... You have to get a lot of young people agitated by the fact that life's going to be worse

with every direction because of the actual debt. And then you've got people like Krugman and Stiglitz saying a 180 degrees opposite, that for a long period of stagnation, you'd need a ton of public investment to create jobs, and make the economy more productive, make up for the short fall in private demand. So, what does a group like yours, when it's looking for best practice, do when it comes across a very important public policy area where expert opinion is completely polarized ideologically in terms of the diagnosis and the remedy?

0:22:48 HK: Sure. So the question is what do we do when we come across an issue where expert opinion is totally polarized? And then the example pointing out with macroeconomic stabilization which is true, that there are some people putting a lot of money and effort into getting our debt down and other people, prominent intellectuals making arguments that we need to be doing deficit spending to help the economy, what do we do about that? It's two pretty different questions, so I'll answer one at a time. One, what do we do when expert opinion is polarized? We have to weigh things like, what do we think is the right side of an issue, but also how confident are we? A lot of these issues I feel like the people who are both relatively informed and analytical, but also humanitarian in their alignment and motivations, they're often all coming down on side of the same issue, where there's not a lot of... The arguments don't seemed balanced to us on a lot of the issues that we're most interested in, and that is a judgment call. And again, Open Philanthropy Project, it's going to be a lot more controversial than GiveWell. That's part of why we're changing the name, and it's going to involve a lot of judgment calls, but I think that's how we're going to do the most good.

0:23:57 HK: But I do think with a lot of these issues, it's like one side has the public interest on its side or the humanitarian interest, which often is different from the American interest, or not exactly identical, and then the other side might have more like trade groups, and groups that represent a certain set of more personal interest or company interest. So, we often see that things aren't polarized. If we feel that they are, and we can't tell which side is stronger in its argument so we're not very confident, that makes us less excited about an issue. That's been a major issue for us on macroeconomics, in general. I do not feel as confident about which side we've picked for some of the grants as I do for a lot of other issues. On the other hand, I do feel convinced enough, I mean I think, again, we could run a whole meeting just on that, given the size of some of the opportunities, and this is where I'm going to say that actually there are...

0:24:50 HK: It's not as simple as the debate between Peterson wanting to cut the debt and Krugman wanting to do deficits spending. There is a lot of other ground in there that I think is pretty neglected and where there is room to do something differently. So, two examples that I think are pretty important. One is neither of those two sides of the argument, although Krugman has opinions on this, are really talking about what the Fed should do. The Fed is different from... The question of, "Is our government doing deficit spending? How much is our government spending?" Is a totally different question from "Where is the Fed putting the interest rate? And more importantly, what happens when the interest rate hits zero and the Fed has to resort to unconventional measures?".

0:25:29 HK: That question, I think, has less of a partisan valence, although it had some. And that's a question on which there's not a lot of advocacy going on. And mostly advocacy that is going on is more on the kind of, "We need to go back to a gold standard. We don't trust the Fed at all" side of things. The idea of someone saying, "Before you decide where to put that interest rate, you should hear from a working mom who cares a lot more about avoiding unemployment than about some of the other factors being considered." That is, I think, a powerful idea. It's something that isn't being

tried. It's something that a group we're supporting is going to try to a small degree, and we're going to see how it goes. And I think there's a lot of potential power there. I also think there is obviously a lot of concerns and your general Open Philanthropy Project grant is going to have much a bigger upside and much bigger downside than your general GiveWell top charity.

0:26:23 HK: Another area that I think is a little bit to the side of the Krugman debate, the Krugman-Peterson debate, something we've thought about is, "Okay, so when in any..." In theory, I think most people agree that the government should be spending more when the economy is weak and less when the economy is strong. But at any given time, the debate is always just breaking down between people who want more government and less government. Could you do a better job designing policies that over the long run are neutral, with respect to the size of government, but are doing a better job making government countercyclical, making government spend when the economy is weak and save when the economy is strong? That, again, if you can design those policies, you may get through a lot of the partisan valence. And you may get rid of a lot of the debates that clutter things up now, where everyone is talking about what to do today, and it just ends up with the left wanting more and the right wanting less. Yeah, Eddie?

0:27:17 Speaker 3: It seems like a bunch of the topics you're going after were politically contentious arguments that... Topics where people have very different views.

0:27:24 HK: Yeah.

0:27:26 S3: It seems like there are other things where everyone agrees that they are good, maybe it's like scientific research fund that's publicly-funded or getting a better computer system for Medicaid and feel like the issue of getting the proper resources to them. And so, you can say agrees this is good, but the government for whatever reason doesn't want to fund it. And if your core thing is like moving money, you can say, "Hey, we're going to be the ones that provide that additional funding to these areas. The government that everyone agrees should be happening, but they just can't get funding they need.

0:27:56 HK: Sure. So, the question is have we... I think the question is have we thought about, instead of going to very politically contentious issues, going in the issues where everyone agrees the government should be doing something better, but they don't have the funding and we could provide the funding. This hasn't been as much of a focus of ours. I think the leverage is just not as great. I mean, I think the kind of the whole point of working in policy is that you can, and there are very impressive cases I believe, where you can spend a relatively small amount of money and you can put a group in position to raise the profile of an issue that people had been ignoring, and then you can have a massive impact. And I think paying ourselves for a new computer system for Medicare has... A, it has some problems because the government has so much more funding than we do, that it's like if you start doing that, you could really wind up just spending all your money on the things that government doesn't want to pay for. And B, there just isn't the same leverage, so that hasn't been our focus. If we ran into some thing that we thought was a really crazy opportunity, the government really ought to be doing something, we should just pay for it ourselves, I think we would but that hasn't been the main focus. Yep?

0:29:06 Speaker 4: You mentioned a couple of times that you had some advances in mind of studies or case studies you did where advocacy was... Those that were cheap and had tons of leverage, and [0:29:15] _____ maybe either one or two... Summarized in a sentence or two. And the second

question is... I think I know why you guys are focusing on the US. I think I've actually asked you this question before. But do you think that in the long run or if you were based in, say Germany, or, I have in mind, India and China, there probably is more leverage available for advocacy in countries like that, sort of bigger and/or kind of third world?

0:29:47 HK: Sure. Two questions. One of them is can I give examples of successful policy advocacy, and two is do I think in the long run that we might have more impact working on other countries besides the US? I'll take them one at a time. One, examples of advocacy being successful. First off, I should just note that we, what we're doing is we're... We've gone through and tried to find all the claims we can, that philanthropy supposedly made a big difference, and we keep a big list of those, and then we take the ones that we're most excited about and we actually have a couple people who we're contracting with, and they're historians, and they go and try to figure out what happened. Then they write up a big report and they talk to as many people as they could who were involved.

0:30:26 HK: This is an example of us trying to do investigation in a way that's not statistical, that's not randomized, that I think is the way you have to do this sort of thing. And I think some of them are pretty cool. The only one that we've put out publicly so far, 'cause these things take a long time, is on the Robert Wood Johnson Foundation's impact on Health Care for the Homeless Program, and just want to remind people, we're looking for cases with these history studies of groups having a big impact on policy for a little money. We're pretty agnostic as to whether we agree with the change of policy, 'cause that's not the question we're asking.

0:31:00 HK: So, that I think, it's on our blog, and if you Google "Health Care for the Homeless Case Study GiveWell", that should be a way to find it, and I think that was... I was like fairly convinced that there was something there. That was largely a case in which by providing a demonstration model of how healthcare for the homeless could be set up, they were in position to then get that added to a bill when homelessness was a big issue, and otherwise it would have been like, they were just... People would have been, "Well we could add healthcare to this bill, we don't really know how to do it, we don't have any examples of it working so we won't do it", and that I think is a common approach in policy, is to try and create a proof of concept, get some ideas ready to go, you'll never know when your political moment will come, and when it comes, it's going to be too late to start to work on them then. Other examples that will be out soon or are out, we have two more case studies coming out. One that I really like, although it could take a little while for it to become public, is on the Affordable Care Act, and the role philanthropy played in that.

0:31:59 HK: Basically the person working on this, Ben Soskis, basically tried to chase down all the different ways in which philanthropy affected it. Some of it was in end game where it was like a big campaign being funded largely by Atlantic Philanthropies. Some of it was way earlier just building up little state groups that created, had a part in creating RomneyCare without which I think Affordable Care couldn't have happened. So, that I think is a really good one. Final thing you can Google on our site, "Center on Budget and Policy Priorities". We had a conversation with them, we posted notes from it, and I think we have not done the historical inquiry, but I am going to guess that those are going to check out and be really impressive. This is a group that basically looks for opportunities to do trustworthy analysis that changes people's minds about what a policy should be, usually in the direction of favoring lower income Americans, and I think they've had some examples of success where just like, based on the story, it's just like, "Yeah, unless this person's just lying that sounds like that was probably a really good one." So those are the examples.

0:33:04 HK: Other question. Other countries, we do work in the US for now, largely for practical reasons, you have to build a network, you have to meet people who are good at this, you have to have a lot of background knowledge. I think that's just infeasible for us to start trying to change policy in France right now or something. In the future, yeah, we've thought about it. We've had conversations, maybe the real place that matters is China, in a lot of ways, China's policy matters more than the US. I don't think there's... There's no country that I would bet right now, is going to turn out, after all the analysis is in, to be more important to work in than the US, or better to work in than the US.

0:33:40 HK: So I think if we looked at all of them, I think some country would be a better place to work, but I thought about the obvious candidates and I think, A, the US is really important, and really huge, and just a world power and what it does, matters, it's a lot, B, the US has this political system that in many ways is more fertile ground for this kind of work than a lot of other systems. It's like a lot more fragmented and there's just a lot more opportunities for interest groups to make a difference, and I think that's a real thing often leads to bad results, but if we're trying to exploit that, it could lead to good results. I don't think that's conclusive, but there's no country I would say right now, "Yeah, this one's definitely better to work in". Yeah, Bill?

0:34:28 Speaker 5: This is more of an observation than a question, but maybe there's a last question embedded in it, but picking up on Bob Kuttner's observation that, particularly, macro-stability issue that say, polarized by ideology. Even though there are very, very technical issues involved, that when I got a chance to talk with you earlier this year and get a little more detail about the advocacy group that was... I think one of the things they went through was Jackson Hole where Janet Yellen was speaking, and they had signs saying, "Want to know about labor dynamics? Ask an unemployed worker". From my perspective, 'cause I know a little bit about this field, I could imagine something with... Anything closer to the zero marginal effect of this moving the needle in that debate. Janet Yellen, as a labor economist, is extremely sensitive to those issues, and so now I'm aware of the arguments of Paul Rubin saying, if you throw in 4% inflation for a few years, you could burn all the debt and all of that, but it just struck me that once you get to those arguments, it's a very political issue. And therefore, I guess my...

0:35:42 HK: Just to clarify, I didn't...

0:35:43 S5: When I heard about that advocacy, that sounded like almost a political action committee that I was affiliating with as opposed to the original GiveWell brand, which was more of the Peter Singer "save a life" type of thing, which conservatives and liberals could agree on is important, kind of independent ideology. But in this particular issue, when I look at Alexander's work, I would say 80% of his sources were left-leaning economists.

0:36:17 HK: On this cause, not on... I would not say that would be true across all causes.

0:36:21 S5: No, I think, in the whole, and that goes to...

0:36:22 HK: Yeah, sure on this cause, yeah.

0:36:24 S5: Oh, right, right. So in other words, my concern... Maybe the question is, is moving through some of these issues actually kind of repositioning you as closer to a political action

committee, which is perfectly fine if you're pivoting to the 0.01, the billionaires who are perfectly free to support whatever causes they would like, as opposed to sort of just to pick one brand. How do you guys think about that kind of conflict in the two brands that are developing here?

0:36:56 HK: Sure. So how do we think about the brand conflict and is there a risk of turning into a political action committee? From a regulatory standpoint, I'm highly confident that everything we've done is totally kosher. We're checking all that out with legal counsel. And there's a lot of space in between distributing bednets and being a PAC. We're not a PAC. We don't support PACs. So, I think there's a lot you can do to influence policy and there's a long list of philanthropy doing it that is not a PAC, and is totally fine from a regulatory standpoint.

0:37:31 HK: From a brand standpoint, yeah, we're definitely changing. But we're trying to split, not pivot, and that's why there's the new name. And that is why we are intending this split into two organizations down the line. We know that these are not the same project. These are two completely different projects, that's what we want that to be. The continuity between them is that what we've always been interested in is, how do you do the most good possible?

0:37:57 HK: I think if we were to say, "How do you do the most good possible?" As long as it's nothing that people disagree with and as long as the right and left can come together equally on it, and as long as it doesn't involve doing things that are controversial and are political, and things that get in the news in kind of a simplified, distorted way, I don't think we're doing the most good possible anymore. I think we're closing off too many things at that point. And I especially think that the... I think it's definitely the case that the most experienced philanthropist would be more adamant than I am that we aren't doing the most good that way. So I'm actually less convinced than they are, but there's certainly a strong perspective out there that says, "These are ridiculous things to close off." You at least have to look into them.

0:38:35 HK: That's what's always driven us as people. That's what we've always been interested in. All else equal, I'd rather work on something that everyone agreed on. I'd rather work on something that wasn't controversial, but all else isn't equal and I think we have to check out all these opportunities to do good. That's what we want to do as people. As far as an organization goes, pragmatically, we're housing Open Philanthropy Project to GiveWell for now and eventually we expect those brands to split.

0:39:00 HK: One more comment just on the effectiveness of advocacy for the Fed. I'm not going to get into exactly why I think what I think what I think. I'm just going to state that we did ask this question very thoughtfully, and put a lot of work into it and talked to as many people as we could reach who would understand how these decisions are made. Alexander did a great deal of research on how the Fed makes its decisions, how the whole structures works, who exactly is being targeted by this group, and why we think they might be influenced. And just to be clear, I don't think the chance of influencing anyone are great. But I definitely don't think it can be ruled out, and I think that what we did is well-considered with respect to that. It's certainly an issue that I was equally nervous about and I feel that we're now on top of.

0:39:46 HK: So, I'm going to go... We only have 15 minutes left and I have two more giant topics to cover, but they will be shorter than that one, I think. I'm going to talk about global catastrophic risks. I mentioned before, these are things that you could imagine sort of derailing or changing the trajectory of civilization. I know that sounds hyperbolic, but I think with the things I'm going to talk

about, it's oftentimes not.

0:40:11 HK: So I guess I'll probably start with climate change 'cause it's a good intro to this space. Climate change is definitely something that could affect the entire planet, could have really magnified impacts. On the other hand, climate change is something that does get a lot of attention from funders, from nonprofits, from advocacy groups. I think there are other causes that are as dangerous or are as important that are getting less attention in general. But I think there's one aspect of climate change that I think is pretty interesting.

0:40:45 HK: Most of the groups out there today are pretty focused on reducing carbon emissions, and that reduces the risks of something really bad happening. But regardless of how much we reduce carbon... I mean not regardless but... If we don't reduce carbon emissions a huge amount, and what looks like a realistic amount at this point, what we had is we have a bunch of economic models that tell us that climate change is going to cause a lot of damage, but basically leave society...

0:41:09 HK: Society's going to be a lot better off a 100 years from now than it is now, even if we don't have climate change. And that is what the projections say. That's what the standard projections are saying. I think the thing that worries me most is that those projections might be wrong. I don't think anyone has a lot of faith in those projections, and I think climate change could be a lot worse than anyone's saying it will be. And the question is, if it happens that it is much worse than anyone is currently projecting, there are very high risk things we could do other than reducing emissions that are a lot cheaper, a lot more reliable, and a lot more effective and also a lot more dangerous to deal with that situation.

0:41:49 HK: So, the classic example here is solar radiation management, which could include something like shooting sulfur into the atmosphere to cool the planet. This is completely crazy in the sense that it's like... It's basically messing with the entire global ecosystem, and we don't know what the effects will be. It may be something that turns out to be something people want to do if climate change turns out to be really, really bad. Another issue with solar radiation management is that anyone can do it, which is actually a really scary thought. Any country can decide that it would like the entire planet to be cooler, and so if any country decides that and there's not a good framework in place, and there isn't good research in place on how to do it as safely as possible, that itself could be a global catastrophic risk.

0:42:33 HK: And so the general area of doing research on geo-engineering and saying, "What do we know about how this would affect the planet? How much could we know? How is... What is the safest way to do this? What is the best way to do this?" and maybe most importantly, "How are we going to decide who's allowed to do this? What's the global governance around this?" That's a pretty empty space and it's been pretty surprising to me that you see the hundreds of millions of dollars going into trying to reduce carbon emissions, but it's like a very sparse area of people trying to research geo-engineering and figure out how we're going to govern it, and to me in some ways this is the most important part of the whole space, because this is the part of the climate change issue that could be most catastrophic, is either needing this technology or using it and having bad results from it.

0:43:21 HK: So something we're looking into and Howie Lempel works on global catastrophic risks. He's currently kind of getting to know the Solar Radiation Management Governance

Initiative. Another one of those names, and they're basically trying to solve that problem of getting people together to talk about who's going to be allowed to do this, how we're going to govern it. I think that's research that ought to be done and we're thinking about whether to support that group.

0:43:47 HK: Other global catastrophic risks. There's the general issue of bio-security, which I think is a lot more selling it today than it was when we started working on it because of Ebola. I think Ebola in the scheme of things is not the scariest pandemic we could have. I think arguably, the flu, if it was the wrong kind of flu, could kill an enormous number of people and be incredibly disruptive. Even worse than the flu could be a hypothetical future disease that comes out of a lab, whether on purpose or by accident that has properties no natural disease has.

0:44:19 HK: How can we be prepared for this? There's a lot of different ways to talk about how to be prepared for this. One idea is to talk about how strong global surveillance systems are. How good are we at catching a pandemic early and responding to it and containing it, rather than having to respond after it gets out of control? Obviously, there's at least some parts of the world where we don't have as good capacity to do that as we could. Strengthening global surveillance could be a way to prevent both the potential flus and the potential synthetic pandemics.

0:44:49 HK: There's a whole bunch of other aspects to bio-security. It's actually... It's a huge space in the sense that there's a lot of government in it and a lot of government money, but it gets almost no attention from philanthropy and so a question we've asked is what does that mean? Do we consider this space crowded or not crowded when there's a ton of government and very little philanthropy. So far our answer is we should probably think of it more as not crowded because there's... The government has certain strengths and weaknesses and so does philanthropy, and we think there's going to be opportunities to do things government can't or won't do or improve the way government is dealing with these things to be more prepared for a pandemic, whether natural or synthetic. That's when we may end up going the kind of like full-time specialist expert in that cause root because that is a very big cause of a lot of stuff going on. By contrast, geo-engineering is a smaller field, but we may go that way with geo-engineering too.

0:45:46 HK: Some other causes, there's geomagnetic storms. So, basically, solar flares, a really bad one could potentially damage a lot of transformers and cause like a really bad blackout that could, according to some people's claims, be long lasting and global, and could be incredibly disruptive. Maybe not as scary as the other two things, but easier to prepare for and we basically been contracting with our favorite evidence reviewer, David Roodman, who's put out a few reports on our site about like, who we should listen to about how scary geomagnetic storms are, and the kind of tentative conclusion is, we not only need more information, we know how to get it, and so there may be good opportunities here to basically learn more about how robust our current system would be to a really bad solar flare. So that's one we're working on.

0:46:37 HK: Other global catastrophic risks. Nukes is one that gets more attention than most of the ones I just mentioned but there may be more to do there. It's definitely an immediate and drastic risk, just the idea of a nuclear war. There's some cross-cutting stuff, like a lot of these risks, including climate change, including nukes, including some other things. The scariest aspect of them would be their impact on the food supply. So are there ways we could be more resilient to a disruption in the food supply?

0:47:06 HK: There's questions around when and how, and what's going to happen when artificial

intelligence gets much more powerful? This is something that, I think, has been in the news somewhat recently, and I think is getting taken a lot more seriously now than it was like a year ago. I think it's something that's like a long way off and it's really unclear what we could do, but basically, our only action on that point right now is Howie is going to go to a conference on this pretty soon and try to, like we think it's a good conference to try and figure out how serious this issue is and whether there's anything that could be meaningfully done today.

0:47:39 HK: I mean I think it's fairly clear that if and when we develop really powerful artificial intelligence, that that will be a situation that has high potential for big positives and big negatives. The question is whether there's anything to do today that would make it more likely to go positive. We really don't know the answer to that. So that's global catastrophic risks. And I'll take questions for 10 more minutes. Yeah?

0:48:07 Speaker 6: What about asteroids?

0:48:08 HK: Asteroids is one we looked into and is not on our priority list. So an asteroid could hit the Earth, and a big enough one could cause a huge amount of death and destruction. So asteroids are pretty well quantified and well understood, and the risks are known to be pretty low. The expected lives lost, in terms of on a yearly basis, is just not very high. And I think the asteroids that are big enough to cause a real global problem, we believe that the money is basically there to track them and most of them have been tracked. I don't think that cause looks as good as the other ones. One of the other general patterns here, is all the things I just named are more human-made problems, that are much more scary as of 200 years ago than they would have been before that. Anything that's been a candidate to wipe out the species for millions of years and hasn't done it is kind of inherently less scary. Yup?

0:49:07 Speaker 7: Do you think the estimates on probabilities for these global catastrophes, can it be inaccurate? Because I'd seen numbers with some of these things, and people who come up with like a 1% chance, 2% chance over some time frame. And it seems like everybody doing research in the space are people who are obsessed with these issues and probably overestimating by orders of magnitude. Is it possible to come up with a reliable estimates of the odds of like AI, or bioterrorism that wipes out millions of people?

0:49:34 HK: Yeah. Is it possible to estimate the odds of some of these global catastrophic risks? And I would say no, it's not possible. With the information we have. I mean some of them, yes. Asteroids, yes. But yeah, for the most part it's not, and we have to live with that. We have thought about what we can do to get the best estimates available. There are processes that are believed to lead to better results than just guessing. And if you pulled together a lot of the most knowledgeable people, and had them go through that process, you may get a better estimate. But it's not really at the top of our list. I think, I think basically with all these... It's, this is such a hard area to compare to others because we are talking about the odds that we actually make a difference are really, really tiny and the potential size of the difference we made is really, really huge.

0:50:24 HK: And we don't know any of those numbers are. And that's why we've kind of thrown global catastrophic risks into their own bucket. We think it's just... It's... We think it's an interesting area. We think there's a lot to be learned. We think it's a good fit for philanthropy, and a lot of them are neglected by philanthropy. So we want to work on them, but yeah, we don't really have a great way of comparing them to other causes and so it's been something where we said, "Look, this one

makes sense, We'd like to work on it and see how good we feel about it in five years". Same with US policy. Same with science. We haven't really had a lot framework for comparing them to each other yet. Yeah?

0:50:58 Speaker 8: What do you think about space research and in particular, colonization or something like that as a solution to global catastrophic risks?

0:51:09 HK: Yeah. What do I think about space colonization as a solution to global catastrophic risks? I think it's, in the long run, maybe the best solution. I think it's a long way away. And also I think... I think there's more funding and support for exploring that kind of thing than for dealing with a lot of the risks that I'm talking about. And it's kind of knowably a long way away. A lot of these other things, they're actually pretty unpredictable. Some of them we know are not going to happen tomorrow. There's almost nothing here that I could really promise you, can't become a risk within the next few decades. And I think space colonization looks pretty unlikely to get there within the next few decades, unless it's assisted by some of these other developments like artificial intelligence. I'm basically out of time to talk about science. I'll just mention that we are also thinking about Science.

[chuckle]

0:52:02 HK: And we have a small team of junior scientific advisors, because I think it's, science is basically, it's like another language, and so you have to have someone to translate even to get at the most basic questions. So you need experts even before you start. I think we do have a lot of interesting stuff we found at science. We're well aware that there's a ton of money that goes into it already. The NIH is tens of billions of dollars a year supporting biomedical research, and so you have to be... The numbers in science, to be a big funder in science are very big, compared to what they are in policy. You can be a major player in policy for much less money. And we know that, and so that creates a challenge. And so our question for science has been, what are the big funders failing to do and why are they failing to do it? In a way that we might be able to fill in the gaps. And the answer to that could be, nothing. But preliminary, I believe, there is quite a bit and we will be writing about it. So that's one of the things I have on my to-do list.

0:53:02 HK: Hopefully, I'll get to it in the next month, just a whole series of posts on what we've done on science because it's been a year since we've written about it. It wasn't a focus of ours this year. It was just something I did when I could when I was working with the junior advisors but we do have some interesting preliminary theories and when we get time, we'll explore them more in depth. I'll take five more minutes of questions or maybe 10, but people can leave because it's 9:30, and then we'll call it. Yep?

0:53:30 S8: So this is going back to policy, what do you think about metapolicy initiatives like Lawrence Lessig's Mayday PAC? Do you there is value in that?

0:53:39 HK: Sure. What do I think of metapolicy initiatives like Lawrence Lessig's Mayday PAC? So one of the causes on our list of policy causes, is like, it's called "improving democracy", which basically means that. We've been looking for an angle that we're excited about, we just haven't found one yet. So one angle is the money in politics thing. It's saying, "We need more regulation to get money to play less of a role in politics." According to our understanding of policy, and according to the people we trust most understanding a policy, which again there's no formula here

except to learn as much as we can to make the best contacts we can that seemed to be experts and learn from them, that one... A, I think it has a fair amount of interest but B, I think it's... There have been campaign finance reforms in the past and I think they've arguably make things worse instead of making them better, and I think it's really hard to design something that I really believe would have a big impact on the problems we see. I just don't... I, and I think generally, the people that we take most seriously on these issues, just don't see that as the main lever by which democracy could get better.

0:54:49 HK: There's another school of thought which you can read about on the Hewlett Foundation's blog, they're trying to become I guess, one of the most transparent foundations out there, and so their blog has definitely got more on it than most foundations. So they write a lot about this. They're trying to think more about polarization, so instead of talking about the role of money, they're talking about the role of how the parties have grown apart over the years, and what we can do to restore kind of centrism and reason to the debate. Also, we don't feel great about that one, but also that's partly because there's a lot of interest in that one, from philanthropy and from foundations. And also, for various reasons, we also don't see that as a great lever either, but you could read about that. That's one of the few cases where our foundation has written a lot of what they're doing and why they're doing it, and how they're thinking about it. Google "Madison initiative, Hewlett", and they've, yeah, there's a lot of stuff to read. Yup.

0:55:45 Speaker 9: So you were just talking about science, I was wondering if you put much thought into policy advocacy based on sciences? Obviously, the government's a huge funder of science, maybe the best way to improve that would be the same, to get them to work on things that they're neglecting or that they think are not important.

0:55:59 HK: Yeah. Have we thought about policy advocacy on science? Yeah, we have and that is actually probably the... It's definitely one of the areas we're most excited about from a science perspective, and reason is that I think the science funders we've talked... The science funders I'm aware of, it's like... I can name a bunch that I think are like, everyone would say they're really, really good on the science, they really know what good science is and they know what the government could be doing better, and they are like they have the right people and the right knowledge to do a great job, those groups tend to be totally uninterested in systemic change in advocacy and... There is a lot of complaints right now and the complaints tend to echo each other. This is actually doesn't... It doesn't feel very polarized to me about what the NIH is doing wrong, and what the tenure system is doing wrong, what the publication system is doing wrong.

0:56:51 HK: And I think most of the science funders that we perceive as best, their solution is to lead by example where most other foundations in other fields would say, for example, "Let's fund some people to make noise about this issue, to come up with specific proposals so we're ready when the political moment comes," etcetera. And then meanwhile, there's maybe one science funder I can think of that's kind of the opposite where I feel they have paid more attention to policy and advocacy than they have to the actual science, and I think they've been surprisingly impactful, but maybe not in a good direction. And so that's an area where we feel like if you could put together the capacity to have really deep knowledge of the science, and have the best scientists on your team, with the capacity to understand how to change policy, that would be pretty exciting. However, those are both capacities we currently lack and so it could be a while before we really get there. Yup.

0:57:52 Speaker 10: How much do you expect to focus on the kind of meta issues in kind of

research when it comes to primary publication while it's entering scientific reporting, and all things like that?

0:58:02 HK: Yeah. How much do we plan to focus on meta issues in science, like publication bias and things like that? So I have been focused on biomedical research just 'cause it's the biggest area and I just wanted to start with something where I felt it would prepare me for other things. And in that area I don't... I think the biggest meta issues have more to do with the tenure system, the publication system, and the NIH system. And I think publication bias... Well, actually, specifically, within the... In medicine, publication bias can be an issue, but in the areas I looked at most which are more basic biology, I think that stuff is actually less of the key issue. So I mean the answer is like for any given area we try to figure out what the most important sub-issues are, and focusing on those, but that's my quick answer to you on that.

[pause]

0:58:59 HK: So yeah. Something that is very different about the Open Philanthropy Project obviously is that it just, it requires a lot more... We can't have everything we say be based on a rigorous study or on some incontrovertible set of facts, and we do make a lot of judgment calls. The thing that we're trying to keep continuous with GiveWell is just our willingness to talk about it, our interest in talking about it, and our basic promise that whenever you ask why we believe something, we try to tell you. A lot of times there's more layers in Open Philanthropy Project, 'cause you need to have people who you believe are experts, and a reason for believing they're experts, and a lot of your knowledge comes from them. But that's the point of continuity we're trying to keep, and basically, where we are with Open Philanthropy is we... I think we have a long way to go. I think all we have right now is ideas we're interested in and we aren't moving a lot of money out the door, and I think there's good reasons for that.

0:59:55 HK: But what I think we're on a path toward doing is basically understanding what issues are most interesting to us, what issues we most want to become experts in, then building the capacity to do that, and down the line, having really great and defensible and well-reasoned giving ideas that we're then pitching to major philanthropists. And I think, we hope to have a major impact with how Good Venture spends its money and also with the whole... When the whole next generation of philanthropists gets as interested in it as Good Ventures is, we hope to have a major impact on them as well. So yeah. It's 9:36, so I'm going to call it now, but I'm happy to stay around and talk after.

[applause]