

**GiveWell NYC Research Event, December 5, 2016 –
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.

0:00:01 Holden Karnofsky: Okay, so yeah, now I'm going to talk about the Open Philanthropy project. Basically the story of the Open Philanthropy project. Elie and I founded GiveWell in 2007 and then in 2011 we met Cari Tuna and Dustin Moskovitz. Dustin is the co-founder of Facebook and Asana and they were asking kind of a similar question to what had made us start GiveWell but very different also. When we started GiveWell we were asking, "Hey, I want to give a few thousand dollars, I have a few hours to think about it. What do I do to do the most good?"

0:00:31 HK: And they were asking more something like, "We are giving away billions of dollars. We have our whole lives to think about this, all the time. What do we do to do the most good?" That is a question we obviously found very interesting and had a similar challenge around it in the sense that if you want to read people arguing about what our public policy should be, what our trade policy should be, there's an infinite amount to read. But if you want to read opinions, arguments, reasoning, about how to do the most good with your money, there's almost nothing and it's very hard to get up to speed, to learn, to hear what the important questions are. And so with Open Phil, like with GiveWell, we wanted to work with Cari and Dustin in partnership.

0:01:17 HK: Their foundation is called Good Ventures and our joint collaboration is called Open Phil. Open Phil is going to become its own organization soon, I think now in a few months but right now, is still a part of GiveWell. But as with GiveWell, we were going to go on a journey of answering that question for ourselves, of figuring out where to give to do the most good, and of writing about what we did as we went so that other people could get up to speed faster in the future and wouldn't have to reinvent the wheel. In that sense, it's like the same mission as GiveWell: Do the most good with your money and help other people understand what you're doing. The way in which it's different is obviously the two kinds of giving, I think, are very different and I think they call for very different approaches. And I think, there's a few differences between giving at a very large level and giving at a smaller level.

0:02:04 HK: One is at the very large level, there's a lot of things you can do that other people can't do. You can create new organizations, you can allow organizations to build in whole new programs or go in whole new directions that they weren't going to otherwise. But I think the most important difference is I just think there's a huge difference between spending all your time thinking about how to give and building a network of trusted staff or advisors versus trying to figure it out in a day, a year, a few hours a year, even a week a year.

0:02:36 HK: That's what I think the really biggest difference is and I think that leads to different fits. I think when you're giving at that larger level, it makes more sense to do things that might look crazy to other people. Do things that are very long term. Do things where you don't have the evidence and you aren't going to know how things are going for a very long time. And so what you need instead is a lot of context, a lot of expertise, a lot of connections so that you can make the kind of softer judgements about what's worth doing and get knowledgeable enough in an area or get enough of a trust network that you can do things that you believe in and others don't.

0:03:10 HK: And that's what I would say the fundamental difference is between Open Phil and GiveWell is that approach, and I think there's pros and cons. So I think people often find GiveWell unsatisfying because they say, "Can't you do more? Can't you take bigger risks? Can't you think more ambitiously?" But when people hear about Open Phil, they often kind of have the opposite reaction, like, "This is too weird. This doesn't make sense. This isn't going to work." And I think it's always good to think about what kind of donor you're trying to be. And I think that different things make sense for people with different kinds of networks and different amounts of time to give.

0:03:43 HK: To talk about what Open Phil does, basically up until 2015 our main goal was to pick focus areas and that was a matter of deciding what issues we were going to work on, what causes, and I think it is very important to have focus areas. I think it's very important to say, "These are the topics I'm going to think about the most and I'm going to become an expert in, that I'm going to know people in, and so that I don't have to rethink everything from the ground up every time I see a new grant proposal."

0:04:14 HK: And I think it's especially important for building relationships that if you have good relationships in a field, you can get the kind of truth that you can't get otherwise. And in certain fields where you don't have evidence that can be very important. And so we try to pick focus areas, and what we did is we looked at a lot of different areas we might work in. So we made a list of all of the US Policy issues where we might try to make policy better and so that can include immigration policy, criminal justice policy, climate policy, etcetera. And we looked at a whole bunch of other

different kinds of causes and we asked, basically, we asked three questions about every cause we considered.

0:04:52 HK: We said, "Is this important, is it neglected, and is it tractable?" So 'important' means we're working on an issue where if we are successful, it will be a huge deal, it will affect a lot of persons, it will affect them a lot. 'Neglected' means it's not already covered, this isn't already flooded in philanthropic money, there's value for us to add here. And then 'tractable' means when we look at the actual activities we would do, we think they have a chance of working, and that's important as well. And so what we did in 2014 is made a large list of all the areas we might specialize in, and we did short investigations, shallow investigations, where we quickly got a rough sense of how important, neglected, and tractable different issues are.

0:05:32 HK: Then for issues that stood out, we went a bit deeper and eventually we chose issues and we said, "These are the issues we're going to focus on," and then we basically spent all of 2015 building staff around the choices we'd made. And this is an example of why it's important, I think, to have focus areas and to be careful about how you choose them because once we decided we wanted to work in criminal justice reform, we were able to hire Chloe Cockburn after a really long search on who wasn't the ACLU, has a background in criminal justice, knows a lot about criminal justice, wouldn't have come to us to work on another issue but for criminal justice can bring a level of expertise in context and connections that we just do not have ourselves. And what we did in 2015 was really try to build that team around those causes. And then in 2016 was the first year that we really felt we do have the team, now we want to focus on actually making grants. Open Phil has been around for several years, and I would say this is the first year that we really focused on making grants. We have made grants in the past. To put some numbers on it, in 2015 total grant making recommended by the Open Philanthropy project was in the range of about \$15 million. That does not include the top charities and I'm not including that in any of this.

0:06:52 HK: And in 2016, I'm expecting it'll be at or near 100 million. There's a big growth there. In terms of the actual areas we work in, just to give you a really quick flavor of what areas we work in and what we've done so far; there's criminal justice reform, we chose that because compared to a lot of other US policy issues, it was unusually tractable. And what I mean by that is that in a lot of political issues, it's very partisan and it's a very big challenge to come in and make a difference. And in criminal justice reform, I think there's some unusual dynamics where there was a big rise in incarceration in the US a couple decades ago, and it was really unprecedented and it left the US incarcerating a massive number of people relative to other countries.

0:07:41 HK: And we've gotten to a point now where crime is way down, we don't believe that's because of the rise in incarceration. But it does make incarceration less attractive to people. State budgets are tight, and you've got this rare opportunity to, in some cases, I would say, bring conservatives and liberals together, and in other cases just find a little bit easier ground than might have existed in the past for reforms. And this is one of the few areas in US policy where it just, it looks like you can help some of the worst off people by cutting the size of the government and that can often be good. Another good thing about criminal justice reform is it's a state and local issue, so you can pick your battles a lot more than if you're just... Something like immigration, it's DC only. Of course, we're also interested in immigration, but I'm keeping things on the shortish side here.

0:08:31 HK: Chloe leads our Criminal Justice Reform work, she also manages Michelle Crentsil who helps her with that work and that was basically Chloe leading the process of hiring her. Chloe is, in terms of our process, Chloe leads our work completely. She talks to everyone, she hears the ideas people are having, she uses her own personal individual judgment to say what is most worth doing and then she writes up the case for a grant she wants to make. We give her guidance on how much total we're looking to grant out in that area. And then we read the case, we talk it over with her, we try to sync up with her. In some cases, and we being the final decision makers on grants are currently myself and Cari Tuna. And in some cases, we get completely convinced and we reach agreement and we make the grant.

0:09:17 HK: And in other cases, we might not be so convinced but we'll make it anyway because of trust, and as long as that ratio is not too far out of whack, we feel good. And basically our first year of Chloe was something around \$10 to \$15 million in grants for Criminal Justice Reform and then for year two it's going to be more than that. And some of the things we've done, we funded the Alliance for Safety and Justice which is the national expansion of the Californians for Safety and Justice group that was behind Prop 47 and later Prop 57 in California, two of the biggest Incarceration Reduction bills that have been passed in a long time. Also, have been supporting a campaign to close Rikers Jail in New York and that's not just about moving people from one jail to another. In order to close that jail, you need to change policies, and you need to have fewer people sitting in jail for less time. And then we've also done some things to support alternatives to incarceration. So longer term more ambitious restorative justice is an approach to handling violent offenses or non-violent offenses, where the person who committed the offense and the victim meet and kind of make a

deal with each other, and that can be an alternative to going to prison.

0:10:37 HK: Another important point on criminal justice is we did have the question, "What is the trade-off?" Obviously, if fewer people are incarcerated, that's fewer people suffering and that's less money being spent by the state, but does that also mean there's less public safety? And David Roodman who I mentioned before spent about a year going through all the literature on the relationship between incarceration and crime. His results are not public yet, but they will be and they're pretty close to final.

0:11:04 HK: And it was really interesting, he basically found all the studies with the best claim to actually looking at whether reducing incarceration would cause an increase in crime. And looking at it in a sophisticated way and trying not to get it confused with correlation. And when he looked at them hard enough and he got the data and he replicated them, a lot of the studies either fell apart or their conclusions reversed. But when you look at the whole thing in aggregate, his overall conclusion was that on the current margin with the way things are in the US today, with the very high levels of incarceration, he would expect a zero impact of getting less incarceration. In other words, if you have fewer people in prison, you will have less suffering, you will have less money, and as far as we can tell, you shouldn't really expect more crime. And so that definitely points toward the kind of reforms we're working toward, and that whole report will be out in all its gory detail.

0:12:00 HK: And then other causes we work in. We work on farm animal welfare. This is a cause... There's a very very large number of animals being treated incredibly cruelly on factory farms, basically, to produce meat, and eggs, and animal products for people to eat. And I think there's a lot of room for debate on how much you should care about the experiences of animals compared to humans. And there's a lot of arguments one can have on this. And actually in a similar style we've been working on a big long, detailed, confusing report about how you should feel about that question. But I will say that if you do take the position, and I think it's an at least reasonable, or defensible, or possible position that you should care about animals some fraction as much as humans. Then I think the things we've been finding in farm animal welfare are just incredibly good deals. So, earlier I talked about top charities, bed nets estimate that you might be able to prevent an infant death for every \$2,000 or \$3,000. With farm animal welfare, the work we've been doing, it might be something more like 50 to a 100 chickens are no longer in a cage for their entire life per one dollar. And that is quite a multiplier and I think if you... Certainly, I value humans more than chickens.

[laughter]

0:13:24 HK: But it's worth asking by how much, and that is a question I feel very confused about. And I'm glad that we are doing this work. I think it's been very exciting. Louis Baller leads our work and he came to us with the insight that there is a lot of room to make a lot of progress on corporate campaigns here. And so, basically, in his opinion if you go to a fast food company or an image conscious grocer and you say you basically are protesting what they're doing and giving them a choice of either having their brand tarnished, or making a pledge to go cage-free, it turns out that they go cage-free, a lot of the time. There's been a huge wave of cage-free pledges. Taking chickens out of cages definitely does not put them in conditions that I feel good about, but it definitely is better than them being in the cages. And basically there's been a wave of pledges over the last year by the groups that we've been... The groups we've been funding have been pushing for those pledges. We believe they've largely caused those pledges, and if all those pledges are adhered to, it's approximately all the eggs in America will be cage-free in 10 or 20 years.

0:14:33 HK: And that's been just massive sized wins and since seeing all those pledges, we've been taking the work international, as well, trying to fund cage-free campaigns in Europe, in Latin America, and also trying to get things started in China. China is where a huge amount of factory farming is going to be, and we've just been trying to build up anyone who seems to be doing useful work there, just to build the field, and just to get more wins there. And I've recently started also looking at meat chickens, boiler chickens, how they're treated, and looking at fish, as well which, should we care about fish? Well, again, there will be a giant confusing report on this soon.

[laughter]

0:15:11 HK: That's our farm animal welfare work. Other work, there's biosecurity and pandemic preparedness. This gets into the topic of global catastrophic risks. But basically we think that philanthropy is uniquely well suited to take on things that are very unlikely but if they happened could derail all of human civilization. Because it's really not clear who else has the incentive to worry about that kind of thing. And so Jaime Yassif and Howie Lempel are our biosecurity and pandemic preparedness team, and their work is newer and a lot of the grants aren't public yet, but basically have been funding, in our opinion, in their opinion, a lot of the best work on thinking through how the world

can be better prepared for the worst case pandemics. Especially, if and when, the state of synthetic biology advances so that by accident or maybe by maliciousness you could get the kind of pandemics that we've never seen in the past. And that's obviously work that's hard to do a cost effectiveness calculation on, but it's work that we've been doing, nonetheless.

0:16:15 HK: And then a final cause that I'll talk about is potential risk for advanced artificial intelligence. This is also in the category. I would say there's two schools of thought in how to help people the most with your money. One is to focus on the people who are the most marginalized, the most suffering, the worst off, and that is a common theme of the GiveWell top charities, of the animal work, of the criminal justice reform in a sense. And another school of thought about how to do the most good is to focus on giant leverage factors that could affect how human civilization progresses far into the future. And pandemics is one thing that could influence that, especially as biology advances. And then another thing is artificial intelligence. I consider AI to be one of the most dynamic and unpredictable areas of science right now. There have been in the last five years or so, problems that seemed impossible and seemed like things that only humans can do, like tagging images, and in some places playing video games just by looking at the screen or learning how to play them without having a lot of data other than what's on the screen.

0:17:22 HK: Those problems have been... We've seen surprising success on them and there's a new set of approaches in AI, where the combination of state of the art hardware and a certain kind of approach called deep learning has produced a lot of really big breakthroughs on previously really difficult problems. And how far is this going to go, we don't really know. It could be that this is just a flash in the pan that improves voice recognition and image recognition, and nothing else really comes of it. It could be that large, large numbers of the percentages of the current economy become automated and that could have unpredictable consequences for the world, and it could be that we get even more extreme effects where I would say that the reason humans have so much more impact on and control over the planet than other animals certainly seems to be intelligence rather than any particular physical ability.

0:18:16 HK: And if there were something that kind of... If there were an AI that were able to outperform humans at some of the highest stakes task such as doing science, I think the consequences of that are extremely unpredictable. And what we've been trying to do with potential risk for advanced AI is encourage the growth of a subfield of artificial intelligence research which is AI safety research. What we've been trying to do is encourage a world where there's a large critical mass of AI researchers that are good both at thinking about how to make algorithms more powerful, but how to make them more reliable, more robust, more stable and how to think through what impacts they might have on society.

0:18:56 HK: Now I think that's the kind of thing that it's a long ways off. It might never matter. Even if it does matter, the work we're doing now could easily just be in the wrong direction or with the wrong emphasis. But if it does matter, could be so world changing that the numbers are competitive with or better than all the other stuff I've talked about tonight. Those are some samples of the work we're doing. We also have supported some groups that are in the effective altruism communities. Those are groups that are doing work around the general idea of trying to spread the concept of doing as much good as you can whether that's with your career or with your money or with other things. And then we've done a series of one-off grants as well when we see a specially good opportunity pop up. Some of those have been on climate or other global catastrophic risks and miscellaneous other things. I had a couple other notes I want to say about Open Phil and then I will pause for questions. Yeah.

0:19:55 HK: We are, like I said, we're spinning out as a separate organization. We're kind of already talking about Open Phil as a separate organization. It's not formally done yet. I think about a year ago, I said I hoped it would a year, but I wasn't sure. Now I'm pretty confident it will be less than a year and I'm probably aiming for something like six months, and we're definitely actively working on that. And in terms of what's going forward for Open Phil, I think that giving went up a lot this year, I don't think the priority next year is to hire a bunch more program officers and expand it even more. I think the priority more is to say, we're facing a lot of decisions we didn't used to face. We're facing a lot of challenges we didn't used to face. We have a lot of questions about how to compare this kind of giving to that kind of giving.

0:20:42 HK: We have a lot of process issues of just wanting a better developed operations team, communications team, logistics team. And I think the focus of the next year is probably to keep giving roughly around the same level, keep staff roughly around the same level, though I think we want to be looking into hiring more operations staff and generally just getting a lot of our... We've got a better sense from what some of the big questions are that we need to struggle with and getting our best thinking on those out there and written down so that we can work from it. That's a lot of the priority that I want for the next year. That's basically the Open Phil story and update and now I'll just take

questions for the rest of the time. Yeah?

0:21:23 S?: In AI, I think there's the whole safety and a side that we're not sure, but there's a side that we're pretty sure that they're going to take jobs away and that will be pretty much manual labor and that will increase inequality. And I think that's a pretty almost for sure assumption that that will happen and it is happening. So what do we do about this increasing inequality that that causes?

0:21:47 HK: Right. What do we do about the possibility, or you're kind of saying it's more than a possibility, that AI will lead to automation of jobs and increasing inequality? Yeah, I tend to agree that I don't know exactly how much automation is coming in the next 20 years or really what the exact consequences will be. Rising inequality seems to me like a pretty good bet as one of the consequences. And yeah, there's a whole bunch of questions you could have with that as a jumping off point. You could have debates over what the best policy response is. Should we have universal basic income? Should we have some alternative to universal basic income, etcetera. We are interested in that stuff. I would say on the AI front, probably the work that I find most relevant to all the other stuff we're doing and relevant to that is this idea of road mapping. So having a better forecast of which jobs are going to be doable by AI and when. I don't think we're ever going to know that, but I think we do a better job guessing than we currently are, and that might inform policy debates and make them more sophisticated. So that's an example of an interest there. Yup?

0:22:56 S?: Are you guys looking into lobbying governments to increase funding to the type of charities that are supported by [0:23:02] ____ and GiveWell? So for example people estimate that \$1 spent in [0:23:04] ____ \$80 to \$130 of local giving and if we really imagine like buying a nice steak to [0:23:11] ____ . So are you guys looking into that side?

0:23:16 HK: Yeah, have we looked into lobbying or just advocating to governments that they do more funding for foreign aid, for the most effective kinds of foreign aid. We have looked into that. That was a cause that we looked into as a potential focus area. I think there is a pretty sizable already advocacy world of groups trying to make the case to the government, that foreign aid and especially health aid is good. So I think it would be a little hard for us to find an incredibly compelling niche there. We have done some funding on this front. So Center for Global Development is a group that is a think tank and they try to come up with ways in which rich country policies could be better and more beneficial to poor countries. We think they do really good work and they were a group that we made an early grant to, to support their work. But it's not a focus area of ours because when we looked at importance, neglectedness, and tractability, we felt that there were other areas that just made more sense for us to focus on.

0:24:15 S?: Did you actually keep track of and document the process, of how you narrowed down the list? So for example if we had a similar idea as you started with a few years ago, we'll be able to look up your reasoning?

0:24:26 HK: Yeah, did we document the process of how we chose? Yeah we did. If you go to openphilanthropy.org, and you go... I think it's under the research top-nav and then you go to 'our process', that's where it's just going to walk through step by step. And don't miss, there's a couple links that go to spreadsheets, and the spreadsheets actually show the causes that we seriously investigated, how we feel about the importance, neglectedness, and tractability. With global catastrophic risks and things about what the worst case scenarios are and stuff like that. So, we didn't document it exhaustively, and one important thing about Open Phil that is different from GiveWell is it's just, it's not as well documented, it's just not as tractable to document it all. GiveWell is really there to serve people who want all the information out on the web where everyone can look at it and critique it even if they're not going to read it themselves. GiveWell has a high ratio of writing to deciding.

0:25:18 HK: Open Phil, we make a larger number of grants and a lot of times it's like, "I believe this person because I've gotten to know them for a long time and had a lot of conversations with them," they believe a lot of other people and documenting the whole thing would just be impossible. So, Open Phil, I think what we do accomplish, what I hope we accomplish is that you can see at a high level how we're thinking, you can see what our process looks like, what criteria we're using, what our basic reasoning is, where you're likely to have a disagreement, what might be worth drilling down on. But we definitely do not have our thinking documented in the exhaustive level of detail that GiveWell has, so that's the difference. Yep?

0:25:56 S?: About a year ago, Open Phil published a report on cultured meat, and I think the conclusion was something like that Open Phil did not see a path to it becoming cost-effective because the media that the meat grows in... Animal free media, it's very hard to do cost effectively, based on the knowledge that we have. So, is Open Phil, are there any updates to that, or is that still the conclusion, or is it...

0:26:32 HK: Yeah, a while ago we looked into cultured meat, which is this idea of using stem cells from animals to basically grow meat in a dish, grow muscles, so that you don't have to... I think if you could do this cheaply, you would have lower carbon emissions, you would have lower environmental footprint, you'd have less animal suffering, you'd have a lot of the abuses of factory farms, which they really are problematic in a number of ways, you could just get rid of. And, so we looked into whether there was promise in funding the science of how to do that. We ended up coming away just not very enthusiastic about that area, for the reasons you said that we just didn't see a path to it becoming cheap. And it's not just the animal free media, there's a lot of things you'd have to do, all of which are just really expensive right now. And so you'd have to find a way to do a lot of different things very cheaply.

0:27:26 HK: I'm not going to say it's impossible, I'm going to say that we didn't really talk to anyone who seemed to have much of a vision of how to get there. We didn't really feel that we had much of a case for tractability at all. We felt that we could do better putting our energy elsewhere, and that includes actually we're more optimistic about people who are making better meat substitutes, out of plants, using more advanced technology. That's an area that has more funding, and that's why we didn't go in there, but we found that stuff more promising. So, the Impossible Foods, they have this impossible burger, it's a veggie burger but they engineer this heme-protein into it, and basically it's bloody, even though it's made from vegetables. And that's a pretty wild and exciting thing, and that looks like it's got a better path to going cheap than this cultured meat stuff.

0:28:18 HK: That's an example of the kind of thing that we looked into that might be important because replacing meat with something synthetic that satisfied people would be a big deal in all the ways I said. It might be neglected. This isn't the cause that you would hear about the most often, and might be tractable. I think tractability is where we have the biggest problem on cultured meat. And we may revisit that because one of the things I haven't talked about is science. Applying our process to scientific research is super hard because there's so many layers of expertise you have to go through. You just have so many things that can be only understood by a small number of people. So, as we get better at our processes, and our team, we may go back to it, but currently, the view is pretty much what it was. Yep?

0:29:01 S?: I had two other ones but a follow-up question for that is, how does that tie in with... I think you guys gave a grant to Good Food Institute which focuses on cultured meat.

0:29:10 HK: So, we did make a grant to The Good Food Institute. The Good Food Institute is a non-profit that is a general purpose promoter of alternatives to meat. Science on how to make a better alternative to meat, anything that can help businesses grow, but they're not cultured meat specific. So Good Food Institute for example, is interested in catalyzing more of this work on highly advanced technology using vegetable-based foods to be more meat-like, and that's something we're more excited about. So, that's the reason I think if Good Food Institute had said to us, "We do cultured meat, that's what we do, that's all we do," I think we would not have been nearly as interested, we were excited about them as a general purpose org. Yeah, follow-up, sure...

0:29:50 S?: So one thing you said, Open Phil is much more happy with high-risk, and is that just high-risk in terms of working or not working, or is it also working or actively doing lots of harm, and how would you take into account really bad outcomes in your [0:30:03] ____?

0:30:04 HK: Sure. When I say 'risk', do I mean just the risk of failing to have impact, or do I mean the risk of doing harm? Are we okay with the risk of doing harm? I would say to a degree we are okay with the risk of doing harm. I think in some ways, almost anything you do and anything you do to try and change the world has a risk of doing harm. The world is just a complicated place. I can tell you stories about how our top charities might do harm. Anything we come up with, I could say how it backfires. I think there's a certain amount of risk. If you want to make the world better, you have to take a certain amount of risk of making it worse, and we're comfortable with that. I do think if we were looking at something where we felt like there was a really big chance that we could really screw things up really badly, and we didn't feel that we had a great purpose on how to weigh that against the upside, I think that's something where it's going to come down to the people involved, and especially to the comfort of the donor. But certainly, I could imagine that just being something where we might not be up for it, but that hasn't usually... Usually, the approach is more to say, do we...

0:31:09 HK: The thing we tried very hard to do, is be as well informed as we can be. To talk to the people who know the most, to have trust networks involving people know the most, and to kind of say, "According to the best available evidence, reasoning, networks, etcetera, yeah there are risks here, but the best available analysis says that the upside is bigger than the risks, and that the risks are largely offset by the good things." Usually when we are able to say that, we

feel comfortable. I can imagine situations where we wouldn't. And, yeah.

0:31:39 S?: Also if I'm... You talked about how you and Cari who will decide if a grant is made or not.

0:31:46 HK: On a day-to-day basis, that's right. Dustin is also the funder, and so he weighs in as well, but yeah.

0:31:53 S?: Okay, cool. And so ultimately there's a lot of influence highly, highly concentrated.

0:31:58 HK: Yeah.

0:31:58 S?: It obviously makes sense, it's Dustin's money, but it could be both. That means there's a high level of pressure on those values being right. If that maybe makes sense.

0:32:08 HK: Yeah.

0:32:08 S?: Could it make more sense to spread out, to let grant makers or let the grant proposals actually just make the promises.

0:32:16 HK: Let the...

0:32:17 S?: I'm sorry. People running your...

0:32:19 HK: Oh, the program officers. Yeah.

0:32:20 S?: [0:32:21] ____ planning.

0:32:21 HK: Yeah, so this comes down to how decision making is made at Open Phil, so it's a little bit complex because I think what we want is we want the people who know the most about a grant to be the ones who are leading all the thinking, leading the strategy, leading the grantmaking. And we know that a lot of times the best stuff, because it's bold, because it's neglected, a lot of times the best stuff will be the stuff that's hardest to explain to people who aren't an expert. And so we have to reconcile that with just feeling comfortable with what happens. A couple of principles we follow, one, just general minimization of decision makers, is definitely a principle we follow. Although we have a lot of money to give away, we're nowhere close to a government. There's lots of other donors who are just as big, and in our view, the world is better off when you have lots of institutions following their own visions and their own values with a lot of integrity and a lot of boldness than if you... At the other extreme, you could have... We could ask everyone in America to vote on everything we do, and I think that would be a different kind of outcome, a different kind of institution. And so that's not who we are.

0:33:31 HK: I think the place of philanthropy in society is actually more to do things others won't, to try things, to be bold, to go with your values, and because of that, we do try to minimize the number of decision makers. Minimizing that, that's partly also a way of empowering program officers. So program officers are the people like Chloe and Lewis and Jaime, who make the decisions about which grants to recommend. And the fewer people that are in the way to block them, the more empowered they are. Now, how exactly do we decide all this stuff? That's where we try to reach a compromise, because I think if we literally just had the program officers make all the decisions, I think that that would work against just the value of donor understanding what's going on and why, it would work against value of us making sure we'd made the right decision of which people to empower, how much to empower them, how much money to put into what they're doing, and basically we would lose track of our ability to make the right high level decisions about who we're hiring, how much money they're getting, and how we adjust that over time, and to lose our knowledge of the situation.

0:34:34 HK: On the other hand, if we made everything make... If everything had to make perfect sense to us, I think that would strip away a lot of the good things that program officers could do, 'cause a lot of times they'll say "X," and we'll say, "Why do you think that?" and they'll say, "I've been in this field for a long time." And do we have to get all the way up to that convinced? I think that often is going to be time prohibitive. So, what we have at the moment is basically a compromise. Sometimes it's called the 50-40-10 rule, although the numbers can vary from program to program. But the goal is, out of the amount of money we're looking to spend, we'd like 50% of the dollars to be things where the decision makers, me and Cari, really get it. It doesn't mean that we saw everything proven the way it was for GiveWell, but it means that we're basically convinced that it's a good grant. We understand the case for the grant, we're onboard.

0:35:27 HK: And then we want another 40% of the dollars, that brings it up to 90 total, to be things where it's like we could see how that could be a good idea. And so there's a lot of grants that I read where I'm like, "I understand how I could think this is good, I don't really get it," and I'm going to say I'll burn a lotta time and say yes. And then the 10% is the discretionary part, and so the 10% is basically, as long as we don't see any show-stopping concerns, as long as we don't see a problem this grant is going to create for us, if the program officer wants to do it, we'll do it. And so that's basically the... And those numbers can vary a bit, but the basic idea there is that I think that gives program officers the ability to do things that we aren't able to see the case for, without making us like totally blind to everything that's going on and then that's just it's... Those numbers don't really come from anywhere principal, they're just number that I like. [chuckle] Yeah?

0:36:22 S?: When you think somewhat longer term, 'cause you said there are other say, extremely large donors, who seem like, "Oh, this is the kind of thing that I want to be doing," how do you think about your relationship with that? Is it just that you put the stuff online or is it more you are creating a general process that other people do? How would they take advantage of that?

0:36:41 HK: How might we connect with other big donors who might want to do similar things in the future? I think, having the website is nice. It's useful. It helps people get a general, initial idea of how we think. It's also the case that there's just a ton of information that doesn't go on our website because there's... A lot of the stuff we do, can be difficult to explain, it can be controversial, it can be sensitive. And often the reason for a grant has to do with opinions of people. The reason for not a grant has to do with opinions of people, and so what we are trying to do is we are trying to share enough info so that people can tell how we're thinking. But there's definitely a lot of info that's not up there. And we don't expect people to read every word, and even if they did, we wouldn't expect them to be totally convinced. So, our strategy for reaching interested big donors is more of the traditional kind of networking and getting to know each other. And that is something that we have put low-hanging fruit. We've had meetings. We've used our connections to spread the word. We've heard from people who are interested. We've talked to them.

0:37:46 HK: I think in the coming year or maybe in the following year, depending on how much we're able to make it work, we may expand that a bit and put a little bit more focused time, because I think for a long time, we were just, we were making priority areas, we were hiring people. We didn't really have much to say about what we're actually doing. I think our product is a lot more interesting now and a lot easier to talk about with donors. So I think we're getting, I think, a better opportunity to start talking to people.

0:38:13 HK: The other thing that I will say is that I think it's definitely in terms of finding very large donors interested in what we do, obviously, there's just a big amount of luck, like, who has that kind of wealth and what are their values. And there's also a big component of timing. People who have already done a lot of philanthropy, they sit on a lot of boards, they have a lot of commitments, they're probably not going to change what they're doing. Whereas people who are too early, they are not thinking about philanthropy, it's just hard to engage. So a lot of times, it's about meeting someone at the right time or being there at the right time. And I think with those in mind, I also think that just the issues are sufficiently deep. And I think in a lot of cases, I expect that a lot of our... I do expect and hope that we will have similar relationships with other major donors, but I also expect that any such relationships in many cases will take years to develop, even after we meet the people and start talking to them and hear their interests. Yep.

0:39:13 S?: What are the one or two program areas, whether for Open Phil or GiveWell, that you would really like to see new candidate organizations in?

0:39:22 HK: What are the program areas I'd like to see new candidate organizations in? Most of our areas, I can tell you some organizations that I'd be happy to see. I think potential risk for advanced AI is definitely an area where we're pretty convinced of the importance of the cause and the issue, but there's not a lot in the way of grantees. It's kind of an issue that a lot of people don't believe is an issue. And so that might be my off-the-cuff response, so just if we could magically have grantees exist, so that we can give them money, that might be where 'cause I think it's a really important cause, and I think that there's a lot of work people could be doing. But all we can do is fund for the most part or that's what we just try to focus our work on. So that's one answer.

0:40:08 HK: I wrote a blog post a while ago, called Science Policy and Infrastructure. I think there's some interesting work at the intersection of science and policy generally, where the best scientists are often not very policy savvy and a lot of the best policy people don't really know... Aren't very deep into science. And so, if you could have policy informed by the best science about how to make science go better and be more useful, I think that I can imagine more

than one organization that could work on that, that doesn't exist that we might want to fund. So those are some answers. Yeah, Josh again.

0:40:43 S?: So I think it was maybe two years ago that you talked about redoing the website, and it seemed to take another year and a half and that there might have been some difficulties there. So, can you just talk to that and if there were lessons learned?

0:41:00 HK: Sure, what took so long with the website?

[laughter]

0:41:02 HK: And what have we learned from that? Let me think. I don't want to say stuff I shouldn't say. So I can check back but first off, it really is a lot of work. There is a lot of content there and getting all that content... So, the experience we had with building a website is we worked with a firm and there was kind of, step one, was getting the look and feel, the skin, the layout, the organization and that just took a ton of our time. And it was a ton of my time and a ton of Elie's time, and a ton of other people's time who are very senior at the organization. Because that website is our public face and so if someone had come along and made us a bright pink website, and said, "This is really pretty," but it didn't communicate what we were, it didn't look sober enough or serious enough, that would be a problem. And so, in order to really be a part of these decisions and make sure that we're being represented in a way we're comfortable with, we had to be really involved in that. And because our time's scarce, it often just took a long time to get through all the decisions it took to make that.

0:42:12 HK: And there's a ton of decisions to make. There's a lot of content on these websites. So GiveWell has its charities and its intervention reports and its pages about the research methodology. And you can click, and click, and click. And to think about how each page should look, how it's going to fit into that, Open Phil doesn't have as much content. And so step two was once we had the basics down, it was just a lot of details about how everything links together and how everything works, and there were a lot of bugs that just had to be slowly discovered and iterated on. And in a lot of cases, the people who were needed to weigh in on something had other priorities. So it wasn't like we were coming back to it every day.

0:42:52 HK: We started this process at the beginning of 2015. Open Phil website went live at the end of 2015. So, that was about a year and that was quicker because there was less content. And it was more urgent. Open Phil didn't have a website. And then GiveWell took two years. That's basically the situation. What did we learn? I've definitely learned lessons about just how to do this process next time, that I don't think are super interesting or important, that I won't go into here. But a big picture thing, I wouldn't say we learned it, but I would say it reinforced it, which is that you can't just outsource outreach. When anything that involves promoting ourselves and going out there and making the case for ourselves, it's just hard to do that without senior staff being involved; it doesn't just cost money, it costs time. It cost more time than we expected. It continues to be the case that we're just cautious about expanding that stuff too aggressively because our bread and butter is research, that's where our focus is. And as long as we aren't short on money, I think that a lot of times that is what we choose to prioritize 'cause it's not cheap or easy to get, even something like a website. It seems like you should be able to pay someone to do it, but I really strongly feel we're not able to do that. Cat, do you have anything to add?

[chuckle]

0:44:09 Cat: So, I worked on the tail end of it. I got involved with the website relaunch in May and we launched in September. So, basically just...

0:44:18 HK: Once she got involved, it was quick, yeah.

[chuckle]

0:44:21 Cat: It's just like everything Holden said about all the little details. We did some amount of content, like rewriting to update processes that have changed, and noticing when things were out of date and flagging that they were out of date. As Holden mentioned, that just turned out to be an enormous project because GiveWell has existed for almost 10 years and has always been a really heavy content producer. So, pretty much the second [0:44:43] ____ it was the big project.

[chuckle]

0:44:46 HK: Also a lot of decisions. There's just like a zillion things on the website where it's like, "We could fix this thing, it will take a long time. Is it worth it?" And those are just decisions that it's hard to have those made at a junior level and so... Yeah. Yep?

0:45:01 S?: So, have you ever looked into infrastructure and how do you think about that? 'Cause I just think it's very interesting. Like I was in Sierra Leone, and it seems a lot of people [0:45:13] ____ where you had to fund raise for.

0:45:14 HK: Yeah. Yeah, how do we think about infrastructure... Do you mean stuff like roads and electricity?

0:45:19 S?: Yeah.

0:45:19 HK: Yeah. Yeah, the bucket I would put that in, I would say just in the terms of the evolution of Open Phil so far, we looked at causes and we chose causes in the policy sphere, kind of advocating for better policy and in the global catastrophic risk sphere because those were two areas where we thought we understood how to pick best causes or how to try to and thought we would learn a lot and didn't already know a lot. And two areas that we've been a lot slower about and are going to be a lot slower coming are scientific research; what lines of scientific research should we be trying to support? Because that's just such a daunting problem in terms of having the right staff and the right expertise. And we finally did hire two senior scientific staff late this summer. They started late this summer. So, I think we're moving on that, it's just that it's a tough one.

0:46:11 HK: And then the other thing is global health and development. Just looking abroad at things we might do to help people in global poverty that are not like GiveWell top charities. Things that are higher risk. That's where I'd put the infrastructure thing. And it hasn't been as high a priority for Open Phil, partly because we do feel like that's an area that we're already doing a fair amount of work in via GiveWell or that we're connected to. And partly just because it's also, there's a lot of depth to it and there's a lot of things to look at and we're not exactly sure how to go about it and we have enough else going on. So, it's for now our priority is really... We have some focus areas, we want to get better at the underlying framework and process and staff behind making grants in them before expanding out a lot more. So, that's a long story as to why. There's kind of these big buckets of causes that we just haven't really gotten all the way to yet and we've kind of looked at them a little bit and had a few conversations but that's about it. Yep?

0:47:07 S?: I'm curious... You just mentioned now two scientific staff. What are their backgrounds?

0:47:12 HK: Sure. So, we recently, earlier this year, I think it was in August or so, maybe July... Chris Somerville, he's a professor at Berkeley, he's on leave right now, and he's a plant biologist, so was very involved in development of Arabidopsis as a plant model. There are certain animals that people study a lot in the lab because they're kind of simple and easy to study and everyone else studies them. This is the thing that people use for simple plants. And he's also done industry, he has quite a broad wealth of experience. And then at his recommendation we also hired Heather Youngs, also comes to us from Berkeley. Basically I hired her because they have a good working relationship, they've worked together for a long time, and she was a strong recommendation. So, in terms of the backgrounds, in terms of our process, as usual we interviewed a lot of people. We looked for certain qualities. Especially people who could help us work through problems we're having and be able to communicate about what we should do in terms that we found convincing and that we could understand. That was a big thing we looked for 'cause that's what these employees do. And then we also hired a more junior scientist, as well so that was a recent thing. And we were kind of doing that search on and off, but mostly on, for about a year before we made any of those hires. Did that answer your question? Cool. Yep.

0:48:37 S?: Can you speak to, at a very base level here, what sets Open Philanthropy apart from a foundation like Gates, or Rockefeller, any of those?

0:48:45 HK: Sure. What sets Open Philanthropy apart from Gates or Rockefeller, or things like those? So there's basically three key distinguishing factors in my opinion. I guess... Okay, so one key difference is that the Gates Foundation's official purpose and reason for existing is to find things that Bill Gates likes. Our purpose is to find things that are the best ways to help people according to our best judgment. Those are maybe a little bit more similar missions than they sound because my judgement and Cari Tuna's judgement are very involved in those decisions, but they are different missions. And our ambition is to be arriving at answers that other people would agree with if they had the same kind of goal of doing the most good and if they reflected sufficiently much on them. And so ultimately when

there's a judgement call about how to do the most good, and it's really unclear how, then sometimes we have to make a guess and sometimes personal values come into that, and ultimately the donor has to be comfortable with everything they do, and that can affect things as well.

0:49:53 HK: But the overall driving mission is we're all trying to do the most good, that's the mission statement. And then that leads to one sense of openness with the word open. That being that we try to choose our causes, choose our focus areas in a way that causes us to do the most good. So your normal foundation, actually the advice we got over and over again when we were getting started was: When you do philanthropy, the first thing you have to do is pick a cause that you're passionate about, that you're personally passionate about, and then you can start doing all the analysis to figure out how to do the best work. And we really have gone the opposite way around, where we've said, the first thing we want to do is investigate causes, not pick a cause until we've learned more and until we've tried to estimate how much good we can do in each one. And so it's that I don't know any other foundation that looked at a lot of causes, rated them on importance, neglectedness, tractability; looked at the broad variety of causes we're looking at and had that kind of framework and tried to do that. So I think in some sense, sometimes people used intuitions where they said, "This seems like a really important problem." And in some cases people made really excellent choices, in my opinion.

0:51:04 HK: But I think it is a different driving mission. And so then, I think that's one real key. And I think that also relates back to the effect of altruism community and its values. I think if you're looking for a large funder that has effective altruist values, do the most good with the resources you have. I do think we fit that bill more than others and we have more of a connection to that community more than others, and I think that makes sense. Then there's the information sharing, and that's related as well. Just wanting people to be able to follow our thought process and our real thought process. So, I think most foundations, most public communications is PR and marketing. And most of our communications we're just trying to explain why we did what we did. So yeah, those are really the key differences in my view. Yep?

0:51:53 S?: What do you see the relationship between GiveWell and Open Phil and individual donors being like in the future? Are you hoping that individual donors will be filling more of the funding gaps in top charities so that Open Phil can focus on what they do best?

0:52:07 HK: Yeah, what do I think is the relationship between individual donors and top charities and Open Phil in the future? I think GiveWell is a really good resource, in my opinion, the best resource for someone who wants to give to charity and does not want to personally become an expert in it, and doesn't have some other expert that they trust in it. And so, GiveWell try to continue to be a service, to people like that. I also think, by the way, having looked at a zillion things, I think the top charities are really good, by any standard, they're really good giving opportunities. Sometimes I think we can do better with the money we're giving away, sometimes I think we can't. But it's never obvious to me and I think the top charities are really good and have... Basically good ventures are still giving large amounts there every year. But I think that primary purpose of GiveWell is to serve that kind of donor. And I think there's another kind of donor that they might be more willing to trust someone. And so I think if you know a person you trust, and they tell you where to give, and you really feel like they know what they're talking about, you may end up acting very differently.

0:53:14 HK: And if that person is at Open Phil, then, or you generally just trust the thought process of Open Phil then you may want to give to the Open Phil a pot of money. And that's something that logistically is not really a thing right now but will be in the future, we have to complete the spin-off first. And then I think there's a third kind of donor that really wants to build up their own expertise, and I think that kind of person we can help them learn, we can help advise them, and that's pretty much all we do. Does that kind of address your question?

0:53:42 S?: Mm-hmm.

0:53:42 HK: Cool. Yep?

0:53:43 S?: So in regard to the last question. So I think Dustin [0:53:49] ____ like nine billion or so and you're giving, what, like 100 million a year, or something? Or, I don't know.

0:53:55 HK: So, I'm not going to talk about their net worth...

0:53:57 S?: Oh, I'm so sorry.

0:53:58 HK: You can see what Google thinks if you want. Yeah. I'm estimating this year's Open Phil giving somewhere

in the ballpark of \$100 million, yeah.

0:54:05 S?: So in the foreseeable future will it ever be the case that if a small donor donates money to Open Phil that's actually going to increase the amount of money that goes out or is it just going to extend the life of the wealth?

0:54:23 HK: Yeah. I think by default if an individual gave to the Open Phil pool of money that effectively you're giving way out in the future when we have spent more time figuring out what we want to do and then we've gotten to the point where we're actually money constrained. Right now, we're trying to move more slowly than that, and deliberately so. I think we're a ways away from that dollar getting spent, if you put it in that pool. I think that's definitely the case. And that would be one thing to consider if you're thinking about giving there. And a good reason not to give there this year, which you can.

[laughter]

0:54:57 HK: Yeah, Jake.

0:54:58 S?: Can you talk about how you set the amount that Good Ventures was going to give to top charities? And then also do you have thoughts on the impact of the elections on US policy rights?

0:55:09 HK: The impact of what?

0:55:11 S?: The impact of the elections on US policy rights.

0:55:15 HK: Great, great. Two of the toughest questions possible. I've got three minutes, alright, let's do this.

[laughter]

0:55:21 HK: So how did we set the amount given to top charities? It was really tough. We spent most of the year trying to ramp up grant making, and when it came time to think about how much to recommend that Good Ventures give to top charities, we're going to have quite a lengthy couple of blog posts on this. It really hit me how many questions play into this. One quick thing to say is, one key question you want to consider is if you give a dollar to AMF or if you don't, you could think of that as either in the First World you give the dollar to AMF, and AMF spends it on bed nets, and the second case what really ends up happening to that dollar is what I was responding to Avi, that dollar gets sort of saved and effectively adds to the pool of money at the end, much later. And so how you compare those two, I call this the 'last dollar question'. In other words, is giving to AMF better than the last dollar Open Phil is going to spend from the pool of currently available capital? If AMF is better than the last dollar, that argues for spending more this year and trying to fill that funding gap and take all the opportunities you're going to have to spend that money. If AMF is worse than the last dollar, that argues for not giving to AMF this year because you can instead save that money. It will do more good later.

0:56:42 HK: I think that a couple caveats to this is that are, one, the last dollar question is totally vexing and impossible to answer with any kind of confidence, and when I did try to go through it, a lot of it was things like, "Well, if I decide that I actually care about chickens a reasonable fraction as much as humans then probably the last dollar is going to be to AMF." Because we'll find things to do that with animals couldn't probably find enough things to do.

0:57:09 HK: But if I don't decide that, now I'm trying to think about what I'm going to think about chickens in 10 years. So lots of questions like that come up. If I write down all the questions I'm facing and I take my best guess at the probabilities of everything and I look through everything, I can come to an answer, but the answer is very shaky. And so then you get to a bunch of other second-order questions such as, if you do think AMF is better than the last dollar, does that mean you want to ramp giving up all the way, give \$200 billion or whatever it is, I actually don't know the total room for funding, and basically take down all the recommendations from the GiveWell website.

[laughter]

0:57:47 HK: And then know that if you changed your mind next year, that you decided you like chickens next year, that you're now going to ramp it down to zero again and all the top charities are going to be laying people off left and right. Do you want to do that? I think not. And also you've got this issue with pushing other donors out which I've talked about in the past which we're not going to feature as much this year. And then if you do decide that AMF is worse than

the last dollar, does that mean you want to take giving to AMF to zero this year? It was a lot last year. That I think creates its own kind of problems and poses its own kind of issues so for a variety of reasons... I partly just value smoothing our thinking a bit and not moving too erratically, especially not reducing budgets too erratically when our judgement is really shaky and can change a lot, and then there's a bunch of other considerations that get kind of gnarly.

0:58:39 HK: That will be up on our blog post. But the bottom line was that my best guess right now is that I think the expected value of what we can do with the last dollar is better than top charities. It's really shaky and I'm really not confident in that. And that called for a lower amount to top charities, but I didn't want it to be too much lower than it was last year. And also a lot of times when we don't know what to do we pick round numbers so that we're not communicating things we don't intend to communicate. So that \$50 million figure compared to \$70 million last year. That's not all the reasoning, that's some of it. If you want to read more, there will be stuff up. Okay, how did the election affect... We purposefully chose cause areas that seemed especially important, neglected and tractable and also we often work on long-time horizons so a lot of times in policy our goal is to put an idea on the map, put it on the agenda, build institutions, build fields, build knowledge, so that if there is a later policy opportunity maybe 10 years from now, you've got everyone ready to take advantage of it.

0:59:51 HK: So because of the long time frames we work on and because we didn't pick the world's most classic, toxic red versus blue issues, because those don't tend to be very neglected and they really don't tend to be very tractable. I feel the set of causes we took is not massively affected here, and it's something we're still thinking through. So it's something, this is not a final answer. But I would say in criminal justice reform, odds look like they've changed for federal reform but that was never our main focus. State and local level it's not really clear that the situation is super duper different. I don't think this really affects biosecurity and Ted never compared with this one way or the other very much. I don't think it affects farm animal welfare very much. I think the implications for artificial intelligence are obvious. No, I don't see a way in which...

[laughter]

1:00:38 HK: I basically think it's unclear and/or no real change for most of the stuff we do, so that's my quick answer, but it's definitely something we're still reflecting on. And there are a lot of little things, things that maybe this was possible before but now it's not because there's a new administration, new set of people or maybe something was not possible and now it is. But at a high level I don't think we need to re-tool everything. And I think, given just like the slow long-term nature of everything we do, you pick causes, then you hire people, then you get to know everyone, then you ramp up your grant making, you can't turn on the dime. I think we'd be screwing up if we were making decisions that had to be totally ripped up and start over every time there's an election that goes differently than someone might expect. All right, taking one last question from Colin. Yeah.

1:01:23 S?: On the election steps. So then you might carry in, I guess some 20 million dollars or a number like that to picking Trump. Picking a side on the red, blue divide. How do you see... Is Open Philanthropy trying to stay separate from that? Is that a commitment to stay away from the bipartisan stuff or...

1:01:51 HK: So Open Phil doesn't do electioneering because 501c3 public charities can't do that so we don't comment on, we don't advise on, we don't give on trying to get one candidate in office rather than another. That's just outside of what we can do as a C3, we're part of GiveWell. That may change in the future because we are spinning off into a different kind of vehicle that does not have that restriction. So that's why we don't get into that.

1:02:21 S?: And also lots of... I would like to say, and I think most people would say, as a parts of an organization, that's a C3 for instance right?

1:02:28 HK: There's a large number of issues where we take a position, and we took the position because we think it's the best thing for humanity, the best thing for people, but the position is clearly identified with the right or the left. And actually we're not always on the right, we're not always on the left. Across our different issues sometimes we're on one side, sometimes we're on the other side. We're not afraid of being partisan in that sense. We're definitely happy to make the case that policy should be a certain way because that's good for people, and if that policy happens to be with what the left wants or what the right hates, or it happens to be what the right wants and the left hates, that's fine, we'll still do it. So we're not afraid of being partisan in that way, but we don't get involved in elections. That's not something we can do. Cool, well, I will stick around to chat for a bit. But again, thanks everyone for coming. It was a really fun conversation, and I hope you guys enjoyed it.
[applause]