

**GiveWell San Francisco Research Event, June 26, 2015 –
Open Philanthropy Project**

This transcript was compiled by an outside contractor, and GiveWell did not review it in full before publishing, so it is possible that parts of the audio were inaccurately transcribed. If you have questions about any part of this transcript, please review the original audio recording that was posted along with these notes.

0:00:00 Holden Karnofsky: So, it's how we did it. I'm recording this, and I'm going to repeat questions just for the recorder, and if anyone wants to take something off or not include it, feel free to tell me at the time, or email me. So, I want to talk about the Open Philanthropy Project. I'm going to be a little bit light on the background and the full story, because I think a lot of you have been here before, so I will give the very basics, but I want to reserve as much time as I can for Q&A, and also for substantive of updates on where we are.

0:00:28 HK: But really quickly, the Open Philanthropy Project started a few years ago, and it was when we encountered Cari Tuna, and Dustin Moskovitz, Good Ventures, and kind of saw that they were working on a problem that was similar to GiveWell, but different, instead of saying, "I have, let's say, a \$100, a \$1000 to give away this year, and I don't have time to do research." It was more like, "I have a very large amount of money to give away, and I have decades to do the work, but still where do I start?" And it was very hard to find a good starting place. We basically were talking with them a lot and felt that it was a really interesting question and a similar situation to GiveWell, where the question of how to give as well as possible is not something that people talk about a lot.

0:01:10 HK: And there's not a lot to read about it, and there's not a lot of information about it, so similar to GiveWell, we said, similar to the top charities project, 'cause it's all at the organization, GiveWell, right now, for the moment. We said, "Let's try to figure out how to give as well as possible," from this perspective, from the major donor long-term perspective, and publish as much as of... As much of what we're finding as we can to help future people learn about this, so that not everyone has to reinvent the wheel. So, that's the similarity.

0:01:42 HK: The differences are pretty extreme, and I'll think you hear them when you start hearing me talk about what we've actually been doing, because GiveWell really emphasizes things that are proven, are shovel ready, you can give now, and you can have a good idea where that money is going and what it's going to do. Open Philanthropy Project, we have... We consider ourselves evidence-based people in general, but that is not really one of the criteria for the Open Philanthropy Project. And instead, we're looking for things that may be very high risk, very long-term, the kind of thing that a big donor is able to work on, partly because a major donor is able to hire staff and really spend a lot of time and really get into issues in a way that an individual isn't.

0:02:25 HK: So, those are some of the similarities and differences. The way that Open Philanthropy works, is that, rather than looking for charities that are proven cost-effective and scalable, what we've done mostly so far, is we've looked for causes that are, what we call, "important, tractable, and uncrowded." So, in other words, we look for problems in the world that we can work on that don't get enough attention for how important they are and for how worth working on they are.

0:02:52 HK: And the reason we look at causes instead of organizations, is because, as a major funder, a lot of times you say, "I'm interested in cause X," and new opportunities come up that you never could've found otherwise, or you can create an organization, or work with an organization to do something differently from what it would've otherwise. So, you don't want to be stuck with the existing set of organizations, and that's why we've taken a cause framework.

0:03:17 HK: Now, what we done is we've... We've basically made a big internal list of causes that we might be interested in. And for the ones that are most promising, sort of subjectively, we've done investigations of varying depth to figure out how important, how tractable, how uncrowded these causes are.

0:03:36 HK: There are... We've kind of divided, just for convenience sake, for keeping ourselves organized, we've divided the things we can work on into four big categories. And first, I'm going to talk about the two that were furthest along on, and I think that will give the best sense for how Open Philanthropy really works, and then I'll take some questions, and then I'll take break, and then I'll talk about some of the other areas we work on.

0:03:58 HK: So, one kind of broad category of cause that we look at, is policy, policy-oriented philanthropy, and this is the idea that, there is a lot of ways, in which, I would say, public policy, whether that's state law, federal law, other things, is suboptimal. And if you could spend money to basically influence policy in one way or another, perhaps by funding think tanks, perhaps by funding research, perhaps by funding public education, and if you then cause policy to change, you can have a very magnified, very leveraged effect, and so, it's kind of a high risk, high reward play.

0:04:38 HK: We have focused on the US for practical reasons, so we wanted to get into policy, but we think it's way more practical to work on US policy than, for example, to work on German policy, even if we thought there was a really important policy change that needed to happen.

0:04:51 HK: In Germany, I think the amount of logistical challenge there would be a little overwhelming, and we've tried to start to with things that... Things that we can realistically do. So, that is one area we've kind of looked at. And within US policy, what we've done is, we've kind of... Again, we've made a list of issues we could work on, and we've asked, "How important is this issue? How much good could we do by a policy change?" We've asked, "How tractable is this issue?" Which in a political context means, "Do we see opportunities to make a difference? Do we see opportunities to actually make some change happen?"

0:05:26 HK: And three, "How crowded is this area? So, are there people already working on this? Who are the philanthropists? What are the gaps?" And we spent most of last year asking those questions and looking into a lot of different things, and near the beginning of this year, kind of a month behind our deadline, our stretched deadline, we pulled together all our information. And after our meeting in DC, with some people that we wanted their input, we made a rank list of the political issues we most want to work on, using those criteria.

0:06:00 HK: And so, one of the issues at the top of the list is criminal justice reform. And we picked that one because the tractability, the political tractability really stands out. So, many political issues, the only way to really get significant change is to get something done at the federal level in Washington DC. Which is not a good place to get things done these days, and especially with the way that the two parties deal with each other, I'm not going to say work together, that wouldn't be very accurate.

0:06:28 HK: Criminal justice is an interesting issue for a couple reasons; one, a lot of the most important areas to work are state and local, so you have the ability to pick your battles to some degree. Two, and I think more importantly, there has been... There is a changing landscape in criminal justice right now, for a few reasons, but basically a couple decades ago tough on crime was the way to be, and since then incarceration rates have skyrocketed.

0:06:56 HK: The cost, the money that states are spending on prison has skyrocketed, state budgets have gotten very tight, and crime has been plummeting. And so, these days you actually have a lot conservatives, the people who used to be pushing the tough on crime line now saying, "Wait a second, we're spending a lot of government money on this. This is actually big government, we would like to maybe cut it." So, you have like Grover Norquist, of the "tax pledge" infamy signing on this website called 'Right on Crime' that basically is trying to pull together conservatives for policy change.

0:07:27 HK: And so, this is a rare opportunity I think to really try to help some of the most disadvantaged Americans by cutting the size of government. And when we look at that landscape, we see opportunities there that we don't see in other causes, and so, that's an example where we figured the importance of this cause is pretty high, not the highest, the political tractability really stands out, and there are other funders, but it's not super crowded and we do see opportunities.

0:07:52 HK: Another example of a political cause that we have prioritized is very different, so I'll talk about labor mobility. So, one of the best anti-poverty interventions we're aware of is when someone steps across the border from a poor country to a rich country. You've the same person, the same skills, often three times, 10 times the annual earnings, and a lot of that money gets sent back home.

0:08:15 HK: And this is a cause where it's kind of the opposite profile from criminal justice, because in terms of tractability, in terms of political opportunity, frankly it's just there's... I can't tell you much positive on this front. There's been a lot of attempts at immigration reform recently, that haven't gone very far, you do have to work in DC, it's politically very tough, but the importance is off the charts. And even a small increase in immigration could be a huge benefit, in terms of humanitarian benefit, in terms of dollars going to globally low-income people.

0:08:45 HK: That's an area in terms of crowdedness, there's a lot of debate over what to... How to handle people who are already here? There's some debate over high-skill workers, and high-skill visas, but there's very little in the way of a voice saying, "We would like more low-skill workers, low-income people to come into the country because we believe that is a huge benefit. And that is an economic efficiency benefit, and it's a benefit to the global poor."

0:09:13 HK: That is an issue, so to compare and contrast the two a little bit, criminal justice is an area where we see

opportunities, and also there's a lot of organizations, and there's a lot of things you could do, and it's kind of a big complicated space. And so, when we looked at that cause we said, "If we were to work on this cause, we'd want to do it by hiring someone to work full-time in criminal justice."

0:09:33 HK: By contrast labor mobility, the part that we're interested in, the part about the low-skill workers almost no one works on it, it's kind of a dead space. And so, we don't really want to try and hire someone because we don't know that there's enough for them to do, and we don't really know who we'd hire with such a thin space. And so, what we've said for that one instead is, "If we saw a great opportunity to make a grant in this area we would, we're not trying to hire someone."

0:10:00 HK: And so, those being two of our top priorities, and we had a bunch of other policy priorities that are ranked in that range or lower, and it's listed in a spreadsheet on our website that's our March update on policy, we came up with the goal earlier this year of making one big bet in policy. And we wanted to make that big bet six months from when we set the goal, so that would be like aiming for August.

0:10:22 HK: And to us a big bet is either hiring someone, or making a big grant, which we were kind of ball-parking at a \$5 million commitment. And what we did is we had a ranked list of priority policy areas to give ourselves the flexibility to say, "We don't need to make the big bet in any one cause." Because you need to get... We want to work on the best causes possible, but you need to get the person right, the organization right. You don't want to hire someone who's not good just because they're the best you could find in your cause of choice.

0:10:52 HK: You don't want to make a bad grant to a bad organization just because it was your only giving opportunity in a certain space. And so, we tried to give ourselves our flexibility, and so, we said, "There's a list of causes we're interested in, our top choice if we could only do one thing would be to make it higher." And our top choice for making it higher is criminal justice, but we would also be happy to make a big grant in labor mobility, or a big grant in land use reform, which is another cause in that category.

0:11:17 HK: And we have hit that goal because a few weeks ago we made an offer to Chloe Cockburn to come on to lead our criminal justice work, she accepted. She's going to start in August, we're very excited about this, and so, now we're moving on to the rest of our list, but that's kind of where we're at. And that's the whole story, and I think gives kind of a picture how Open Philanthropy works. That we looked into everything we might work on, we looked into some things a little bit, some things a lot.

0:11:43 HK: We ranked things, we gave ourselves options, and then we set a goal and now we've made a big bet in a sense in the next phase is when Chloe comes on and starts basically putting... Making grant recommendations, where we're originally aiming for \$5 million a year of grants in criminal justice reform. I'll talk about global catastrophic risks and then I'll pause for questions.

0:12:06 HK: Global catastrophic risks have similar story, so the idea here is that we are interested in working on things that could... Risks that are low probability, but that you could imagine spiraling out of control to the degree that they affect civilization on a global scale. This is another opportunity for leverage, in a sense, because there is no particular actor, no government, no company that really has the incentive to worry about the biggest form of the risks, the things that threaten all civilization. And if you believe in certain kinds of power law dynamics and you believe that those risks are not as unlikely as they are terrible, then you're just dealing with really big issues that you may be well positioned to address.

0:12:49 HK: Last year, again, we made a list of the things that we thought could really disrupt civilization on a global scale. We ranked each one on importance, which in this case sort of means scariness. It means, what are the scenarios, the most likely scenarios for a major civilizational disruption and how likely are they? So we ranked them on importance/scariness, tractability... In other words, what could we do to prepare for this risk today and, crowdedness, who works on them.

0:13:16 HK: Top priority, as of the time we announced our priorities, was bio-security. So the idea there is that... There, one of the biggest threats, I think. If you told me that something catastrophic happened to all of civilization in the next 100 years, that was really disruptive or even made humanity go extinct or something, I would... One of my top guesses would be a pandemic. Natural pandemics.

0:13:39 HK: So there was a flu shortly after World War I that I believe killed more people than World War I. That kind

of thing is hard to predict. That kind of flu could happen today and it could potentially be more damaging 'cause we have a more inter-connected world, it's kind of hard to say. There's considerations both ways. Also, if you're thinking about something that would have the highest potential for real extinction, you would maybe think about synthetic pandemics. You would think about what if biology advances to the point where we can actually modify viruses to have really bad properties such as being contagious for a long time without symptoms before they manifest, and so, getting spread really across the globe.

0:14:19 HK: So this is, I think, the importance, the scariness is high. There's a lot you can do to prevent and to better respond to pandemics, both natural and synthetic. And in terms of crowdedness, it's certainly something, there's public health attention, there's government attention. But there's very little philanthropy, and we thought there might be opportunities. And after kind of looking at the space and landscaping and talking to a lot of people, we thought that there were a decent number of opportunities.

0:14:45 HK: So we made that our top priority and that scenario, a big complicated space, a lot of stuff to do, a lot of stuff to digest. We'd want to hire someone. Our second priority or kind of our first at the moment is artificial intelligence risk. This is the idea that a lot... There are, I think, some good reasons to believe that we could see extremely advanced artificial intelligence this century. You could even imagine something that was sort of better than the best humans, that pretty much everything humans can do with their brains. And that's somewhat of a very exciting thought, somewhat of a terrifying thought, depending on your perspective.

0:15:24 HK: But one thing that does seem true is that the field of artificial intelligence is advancing really rapidly to the point where this is... This is not just a sci-fi thing, this is something that I think really could happen in the next 100 years, really could happen in the next 50 years. And in terms of whether... I think one of our big questions is whether there's anything we could do today to make that a safer transition when it eventually does happen. And when we set our priorities, we kind of put it on hold because we knew that there was a conference coming up which happened in January that was really pulling together a lot of top computer scientists to ask, "Is this a real risk? Is this coming soon enough that we should worry or could it be? And are there things we could do today to make it safer?"

0:16:07 HK: And an open letter came out of the conference sort of saying, "Yes, there's... There are real issues here," according to pretty much any reasonable definition, I think, of the top people who would know, which I think it's pretty questionable, who would be qualified to speak on this, but I think the people you would name; top computer scientists, people who are maybe more futurist, they were there.

0:16:29 HK: And that open letter, looking at who signed it, I think best guess says, "Yes, this is a real issue. There are things we could do today that could lessen the risk." And that was accompanied by a \$10 million gift from Elon Musk to fund basically a sort of things, but including a research proposal. Sorry, a request for proposals. A request for proposals for people to do research on how to make artificial intelligence safer if and when it does come.

0:16:57 HK: And so, that was something we had been kind of watching and waiting, and we hadn't prioritized the cause 'cause we wanted to see how that conference would go. We did send someone to that conference and after it came out, we kind of said, "Yes. This is a real issue and we see a giving opportunity here because we aren't sure there's going to be enough money to fund sort of the creation of this new field of artificial intelligence safety research." So we've been monitoring that very closely and that's actually ended up being kind of our top priority over the last couple months in this area, is monitoring that situation. We actually... We're kind of in communications now about the late stages of reviewing the research proposals and asking, "Is there enough good stuff here and is there enough money to cover it or should we put something in?" So that's been a major thing for us.

0:17:43 HK: So again, same basic template, we looked at a lot of things, we looked at them on a bunch of factors, we ranked them, we gave ourselves options, and then we prioritized this year. And I think one of the big changes with Open Philanthropy that you're hearing is, we used to talk more about processes and procedures and principles and having lots of possibilities, and in these two areas this year, I mean we've mostly thought about like a couple things at a time, so we spent a lot of time on criminal justice trying to make a hire. A lot of time on AI thinking about this grant. Sometime on bio security thinking about hiring, and we've gotten more focused, but I think focused in a way that comes from having considered as many possibilities as we could.

0:18:20 HK: So I'll pause there, I'll take questions and then at some point I'll interrupt the questions and talk about science. Yep.

0:18:27 Speaker 2: So I imagine the policy area, tractability and crowdedness changes a lot and frequently, how frequently are you relooking at those issues?

0:18:36 HK: Sure. So how frequently do we reassess importance, tractability and crowdedness? I mean it's something we don't know yet. One thing I'll say is that we, you know, we spent last year mostly doing this, and I think this year I've been quite comfortable to focus more on execution. Focus more on hiring, grant making, and not be worrying that our research is already out of date, and as we've kind of... The more we work on hiring, and grant making the deeper we get into a cause, so we can at least see that cause. We can see if it's like no longer important, tractable and uncrowded.

0:19:11 HK: I expect these reviews we've done to not get horribly out of date very fast. I think it's... We're probably talking on a scale of several years at least. That's just my feeling. You can look at the research we've done, you can ask how fast that stuff would change, most of it doesn't seem like it would change tremendously fast.

0:19:28 HK: One clarification I would definitely make is that there are different definitions of tractability for politics, so one definition is like there's a bill on the House floor right now. There's a media story going on right now, this thing is hot right now. That is not what we mean by tractability. I think when we're at that level of action I call that too late for philanthropy. I think if there is a bill on the House floor, and you're just starting to learn about this issue think about who you want to hire, think about what kind of organizations you want to support. No, by the time you get your, by the time you get it all together I think that moment is going to have passed, so that's not what we mean by tractability, and that's why in large degree we haven't used the concept of tractability a huge amount in policy, because most policy issues you look at them and you go politics is unpredictable.

0:20:18 HK: An opportunity may arise, it may not, it's really hard to say, but criminal justice I think is an exception, and I would say there is the rare exception, marriage equality, marijuana policy and criminal justice are the three things I'd point to most for sure where I feel like there's something durable going on where I really expect the next five years to be consistently a better opportunity for making change than like an average year and an average cause. Yep.

0:20:44 Speaker 3: How much interest has GiveWell gotten from and tried to solicit from wealthy individuals in working on these issues?

0:20:52 HK: Sure. How much interest have we gotten from wealthy individuals working on these issues? In terms of the component of our strategy that's about sort of reaching donors and influencing them, our main priority so far has really been think about helping Good Ventures determine where to spend its money, which is a great deal of money and we're no where near having enough recommendations for all the giving they want to do. Think about, focus on helping them and more importantly focus on what that really is, is building our knowledge base, building our research, building our team, all the stuff that I'm describing about doing the investigations, having our priorities, having our team, having our grants, that's really our priority and it's similar to the strategy GiveWell followed where what we really are all about is getting things right more than we are about marketing, and we figure that when we know more, when we have more to offer the marketing will come easier, so I think that's something that actually did work pretty well for GiveWell.

0:21:52 HK: Like for GiveWell we tried marketing early on it didn't really go anywhere, then our product got better and kind of a lot of the marketing really did take care of itself, it's not that we did nothing, we put in effort, but we really picked low hanging fruit and we didn't put in massive amounts of effort into marketing, so I think it's a similar story for Open Philanthropy. It's like we really right now are just trying to build something that's good, and then we're going to focus on selling it. That said, we do some low level networking, and we do try and talk to people who we think are going to be future philanthropists, and I'm optimistic about the long term future.

0:22:26 HK: It's a little hard to say because I think if you talk to someone about philanthropy before they've thought about it it's not going to go very far. If you talk to them about it after they've already formed a lot of relationships, they sit on a lot of boards, they make a lot of grants, you're probably too late, and so, I think a lot of what this is about is someone finding you when they're just starting to think about philanthropy. That's the ideal. That's how we met Carrie.

0:22:51 HK: And so, for me a lot of it is being around, networking with people before they've thought about it, so we're there when they do start, but this year we did announce a new partnership with Mike Krieger and Kaitlyn Trigger. Mike being co-founder of Instagram and that is another partnership where Kaitlyn is coming into the office a couple days a week, they've made a public financial commitment to the grants we're going to make, and I think that's a similar situation like we met them just as they were starting to think about philanthropy, so I think that's a positive sign for the future. It's not our main focus, but I think it's something that we see reason to think will happen. Yeah, Jeya.

0:23:26 Speaker 4: Is there any cause in either US policy or global catastrophic risk that you think is like really important maybe up there with the things you've prioritized, but that you've deprioritized because of tractability?

0:23:38 HK: Yeah. Is there any cause that's really important, but that we deprioritized for tractability? There's certainly causes that are high importance that we de-prioritize because of crowdedness, so you know, climate change, I think both on the po... You could call it a policy issue, you could call it a global catastrophic risk, I think the case for importance is stronger if you're thinking of it as a policy issue, if you're focusing on the total cost to humanity and not about the disrupting civilization risk and that's something where I think it's just as important as the cause we've marked important, but it gets way more attention than those causes.

0:24:14 HK: So there's certainly stuff like that. I mean tractability has not been of huge thing with us for politics 'cause I think tractability's really hard to predict. And so, I think labor mobility, immigration is about as intractable as it gets and we're still interested in that one, politically. And then on the global catastrophic risk side, we certainly were not sure where AI was going to end up on tractability 'cause you're trying to do research now to prepare for some technology that doesn't exist yet, may not for decades or more.

0:24:45 HK: So that could've gone that way. Certainly super volcanoes is something we looked at and said, "We have no idea what to do and that's a factor." Though the probability is on super volcanoes being super disruptive is we don't think as high as the other stuff I named, so that's a partial answer.

0:25:05 S?: Can you give us little more insight on what is Chloe going to do after she starts?

0:25:09 HK: What is Chloe, the criminal justice program officer, going to do after she starts? So, partly I'm going to say, I think she is going to lead the way on what she wants to do. I mean I think we did a really serious job search and thought really hard about who we were comfortable working with, we didn't have to hire anyone, and I feel that I have a good understanding of what she wants to accomplish and I have a good sense that she really knows the field a lot better than we do.

0:25:41 HK: And so, I think in some ways it's going to be like I don't want to say like, "This is Chloe's assignment" 'cause I think Chloe is going to lead the way and I think if things change in the next couple of months for example, and there's something that she says, "This is the best thing to work on." I mean that's very likely that that's going to be the thing she works on. Obviously, it's a conversation, what I'm trying to say is that things that happen in the next couple of months and what Chloe comes in wanting to do is going to be a major factor in what she works on.

0:26:09 HK: But obviously, I think one of the first things we're going to do is we're going to run through a lot of our hypotheses about where the best general areas to be in criminal justice are and evaluate them and pick the ones we most want to prioritize, and so, some examples of that, I think one of the things we're interested in doing is working toward, working toward, meaning in support of not necessarily always directly on, campaigns or valid initiatives that are really seeking a direct sense in reforms sort of proposition 47 that passed recently in California reclassifying a bunch of felonies as misdemeanors, I think that was like a really good thing, and arguably there's political appetite and clinical opportunity for more of it.

0:26:52 HK: Another possibility is thinking through prosecutors as an interest group, as the interest group that often blocks criminal justice reform like a lot of times these days you look at some of these reform efforts and it's like the left is on the reduce incarceration side, the right is on the reduce incarceration side. Who's blocking reform? It's the prosecutors. And so, thinking about why do the prosecutors behave that way, what are their incentives, and how can that be changed and is there a way to have a vision of a good prosecutor as someone who trades off cost and benefit well as opposed to I think the stereotype of a good prosecutor being someone who locks a lot of people up.

0:27:35 HK: That's another area of interest, so I think I could probably name five ideas like that and there will probably be a few that we're most intense about and really looking for grantees on those areas and there're also some grantees like there are some potential grantees that we already have in mind, that we think we need to see their plans, we need to talk to them and check everything out, but we already know who they are and we can already imagine giving substantial gifts.

0:28:00 S?: Part of your thinking in like doing a hire was things like you don't like three strikes laws or mandatory minimum or things like that so targeting those kinds of things...

0:28:07 HK: Yeah. So I described political interventions like campaigns and prosecutors, in terms of the policies we're seeking, yeah, our, so I'll say our general sense incarceration has skyrocketed in the US since the '60s. It's really high relative to our, to what it's been historically, it's really high relative to other developed countries and our basic impression is that it could come down a lot without reducing public safety maybe while increasing public safety and especially penalties for non violent crimes, especially drug crimes, yes, I think getting rid of a lot of mandatory minimums, some other things like when you're on parole or probation and you have a technical violation meaning you didn't show up for an appointment or something, does the penalty for that have to be so draconian? Is it going in and serving the rest of your sentence.

0:29:02 HK: Bail reforms, is money bail a good system? Is that actually something that having people who have more money get out and people who have less money sit in jail until their trial, is that something that's doing society a lot of good. I think these are, our intuition is that, in general most of the opportunity for improvement is in terms of reducing incarceration.

0:29:25 HK: That said, we are kicking off a literature review right now to make sure we've really covered our basis in understanding what is the relationship between incarceration and crime? How much benefit do you get for different kinds of incarceration and is that enough to make up for the cost? The cost, the fiscal cost, the human cost of incarceration, and so, I think to that, that will hopefully deepen our understanding of where the best policy changes are, but I think there is a set of policy changes that we think are pretty likely to come out looking really good and likely enough that we are ready to make that hire. Because we certainly made the hire, I mean, we made the hire knowing that the person we're hiring is looking mostly to reduce incarceration. It's hiring a person, it's also hiring a set of beliefs and biases. Yeah, in the back.

0:30:19 S?: So you've said in the past that a good researcher would be worth more for you than a million dollars? As a person who would be interested in helping GiveWell to do its mission, what is the trade-off that you would want it to make, time-adjust it, between money now, money later... Large quantities of the money with low probability, small quantities of money with high probability. Good researchers, good grantees and people not wasting your time by proposing bad grantees and bad researchers? Like, if there was a well-defined set of requests for assistance, tractability value, probability tradeoffs from people helping you, that would be very helpful.

0:31:09 HK: I think what you're asking is kind of like what do we most want?

0:31:13 S?: Yes.

0:31:13 HK: Do we most want more people...

[laughter]

0:31:16 HK: More big dollar-giving opportunities? I mean, it fluctuates a little bit. It's somewhat dynamic and it changes with how we are 'cause I think a lot of times, the way things work here is that we... A lot of times, what's holding us back from doing more is we don't have enough people. But what does that really mean? You can't just throw people in. It's like, everyone who comes here, we need to train them and get to know them and understand what they can do. And then they need to develop and that takes our time. And so, hiring is not something where you just increase the size of the team. It's more like, I invest my time in the hopes of eventually getting better. So a lot of times, it's like, one of our bottlenecks is we don't have enough people, but a meta-bottleneck is that we don't have enough people to train and manage people.

0:32:00 HK: In terms of what do we most want? I can say right now in kind of a narrow sense what we most want. We are trying to hire. That is one of the main things we are trying to do on the Open Philanthropy side. GiveWell has done a lot of hiring recently, and so, this will be a somewhat different question for GiveWell. But on the Open Philanthropy side, there are still some positions we're trying to fill. There's a Life Sciences advisor role that I'll talk about in a second. We would love to hire maybe one or two more program officers for Bio-Security and perhaps, factory farming or macroeconomic policy. I think after one or two more, we would definitely pause because we need to see how that goes and learn about it. So those are things we want.

0:32:44 HK: I would say secondarily to that probably, we want giving opportunities. So a great giving opportunity in labor mobility would be really, really exciting. We want contacts, so we want people who can help us navigate these

fields. But yeah, actually, I would say giving opportunities in hires are definitely at the top of the list. Those are what our goals are and that's what we measure ourselves by. And I think, wanting more donors is also a thing that we want, but it is not as high up as those things.

0:33:15 S?: So, I can work for factory if I'm in bioscience [0:33:17] ____ advisors?

0:33:19 HK: Sure. Yeah.

0:33:21 S?: Once you've identified the policy area, what sort of levers do you see yourself using to make change? I can imagine talking to politicians, I can imagine public advocacy, I can imagine funding research into something. How visible do you want to be in terms of the Open Philanthropy? Is this our agenda or is it more a behind the scenes thing?

0:33:39 HK: Sure. What are the levers for influencing policy? We wrote a blog post about this. That was pretty basic, but I think it's pretty helpful called "The Role of Philanthropic Funding in Politics" or in Policy. We listed a bunch of things. I mean, there's a lot of different things you can fund if you want to influence policy over the long run. I think, a lot of people, when they think about funding and policy, they think about paying for campaigns. I mean, [A], that's not something we can do as a public charity, and [B], I don't think that's always the most important or effective way to do things.

0:34:16 HK: Some of the things you can do... I mean, I think certainly research can change the way people think about an issue or the way they approach it. Aside from academic research, policy analysis. Sometimes, what you need to do is take an issue and hammer out. Alright, if we had the political window, what would a good bill look like? What would a bill look like that would actually work and that we could get political support for conceivably if the moment arose?

0:34:40 HK: Another thing you can do is, you can build groups that bring pressure or bring a voice. That could be grassroots organizing which is basically building lists of people who care about an issue and getting good at mobilizing them. That could be sort of more of what people call grass tops where you're looking for allies in DC. One group we fund works on immigration and is interested in trying to pull together small businesses and get them signed on to an agenda of higher immigration.

0:35:12 HK: Those are definitely some of the main examples. I think, we try to be just really open about it. We try and pick our cause, but then we're really open about it. One of the things that we looked for in our search for a criminal justice program officer was someone who we felt... And this is a challenge. This is not like an easy... This is not something that you get this experience in a lot of jobs. It's going to be hard to find this person, but we were looking for someone who we felt could consider a lot of these different levers and choose between them intelligently. So there's a lot of options.

0:35:41 HK: I'm going to pause. I'm going to talk about science for a second and about our other stuff real quickly and then I'll take questions for the rest of the time.

0:35:49 HK: I've given a sense for where we are in policy, given a sense of where we are in GCRs, global catastrophic risks. Another area of philanthropy is scientific research funding. The case for leverage here is that if you can fund basically an innovative new insight, that insight becomes kind of a global public good, anyone can use it. And so, your impact is kind of unlimited, and so, you could imagine leverage there. This has been a really challenging... This kind of category of philanthropy I would say is kind of the size of all the others ones combined, just in some vague sense of, "How much stuff is there to look into? How much work would it take to really understand all our options?"

0:36:31 HK: And so, we're taking a much more shortcut approach. I think it may... What we do is going to end up being a lot less thorough, a lot less comprehensive than what we did for policy, and it could take us longer too. But, you know, in terms of preliminary, what we're looking at kind of broken the scientific world into what I call neglected goals and systemic issues. So neglected goals are things that you could do scientific research on, but that maybe don't get enough attention because the goals are not deemed important enough.

0:37:02 HK: So for example, cancer, a lot of people work on cancer, not as many people work on malaria. This is arguably, I think probably because cancer affects many more wealthy people than malaria does and malaria affects mostly people who are very low income. So that's a classic potential neglected goal. You can imagine that what you want to do is you want to direct money into malaria because the existing system is not going to do it enough. This is actually something that I think some other funders like the Gates Foundation believe. And so, whether it's still true with

them around is definitely an open question. But that's one category. And so, we have a list of neglected goals, we're not going to be able to look into all the good ones but we're going to try to talk to broad people to figure out which ones are most worth looking into, and we have some investigations ongoing. And an important part of those investigations is science advisors because we don't have, we don't have the ability to get all the way up to speed on all the biology we would need to know. For example, to understand giving opportunities in Malaria.

0:38:01 HK: So we do have right now a couple of, I think, very talented science PhD's and [0:38:09] ____ Lilly Kim and Melanie Smith working with us and we're trying to hire one or two more full-time people for that role. And basically, they help us understand the giving landscape and the opportunities. And then, another area of science is systemic issues. So this is things where people say, "The system that we have for funding science is dysfunctional in a way. It's not that people don't care enough about the end goal. It's that they're not going to accomplish the end goal as well as they could because of the way the tenure system works, because of the way the publication system works, because of the way the NIH works, because of the way academics incentives work." And so, there, it's basically having conversations with scientists trying to understand how their incentives might be distorted and thinking about what we want to do about it.

0:38:56 HK: So we wrote a post called 'Science Policy and Infrastructure' that is about this issue for Life Sciences and kind of puts forth the idea that actually one of the things that might do the most good when it comes to Biology and Life Sciences is actually some kind of advocacy, some kind of political work to try to modify the way that the current system rewards and incentivizes scientists.

0:39:20 HK: So those are the things we're working on, we're trying to hire science generalists. We have that posting. I believe that posting's on our website or if someone emails me I'll send it to them. And there's also social sciences and there's a whole bunch of fields. And how we're going to get through all of them? We're not going to look at everything. We're not going to do a shallow investigation of all the good stuff, but we're going to sort of do as much as we can for about a year and then take our best guess. And then, we'll have a set of priorities that looks hopefully something like the policy priorities and then we'll be looking to hire and make grants.

0:39:52 HK: Other stuff that's going on in Open Philanthropy in general, a fourth category is global health and development. We already have the GiveWell top charities, but there are other things you can imagine doing that are focused on developing world aid that don't follow the GiveWell model. So for example, "Could we do good by getting the AID community as a whole to pay more attention to mental health?" Some people have said that to us. I'm not sure whether it's true. Could you do a lot of good by trying to strengthen the general capacity of health systems?

0:40:23 HK: This is not something that's going to win on the GiveWell top charities model, but it might be a good idea anyway. Kind of a high-risk, high reward thing. That is something we're not working on right now. We just don't have the capacity. We'll work on it when we do have the capacity, when we've gotten through some other stuff. I don't feel terribly impatient because I have enough familiarity with those issues to feel like GiveWell top charities are actually reasonably high up there on the best things you could fund in that area, and so, we are prioritizing, getting good at things that we know nothing about.

0:40:51 HK: And then, just some other piece of information... We have a History of Philanthropy project where we try to check out things that we've heard about where philanthropy has made a big impact. And it's actually historians working on this. We've published a couple of case studies. And the goal of that is to just try and understand the backdrop better, make us better informed about what has worked, what can work.

0:41:12 HK: And then, another very important piece of information which I think I've said before, Open Philanthropy Project, our vision is for it to become a different organization and we're hoping that happens in the next year. We're not in a rush. It can be a pain to separate organizations. But we think it's important. And we're launching... We're doing a website redesign soon where we're going to have a separate GiveWell website, Open Philanthropy website.

0:41:34 HK: We think this is important because otherwise we talk to people about Open Philanthropy and they start off with a lecture about how not everything is proven as bed nets, and so, we can reduce confusion there. And then, on the GiveWell side, I think we don't necessarily want our views on mass incarceration and labor mobility to be associated with GiveWell, which I think GiveWell top charities, which I think is a lot of what it brings to the table is the credibility and neutrality and the fact that there's not so much controversiality there.

0:42:04 HK: So that's just something we're headed toward. We're not in a huge rush. We're not feeling a huge pinch, but

it's something we're planning. So with that, I'll go back to questions for another 15 minutes, and you don't have to ask questions about what I just said. If you still have questions about policy or anything else or GiveWell or whatever, just go ahead. Yeah?

0:42:22 S?: I thought that I had seen a hire, a open position for around factory farming, but you didn't really address it, is that still happening? Or what is the status of that?

0:42:32 HK: Sure. Are we still looking for a factory farming hire? Yes, we are still looking for a factory farming hire. I don't know. I'm trying to keep my remarks like contained enough. That's another...

0:42:43 S?: Is that one... Is it one of the priority areas in the policy?

0:42:46 HK: Yes. Yeah. So it's like basically there are three areas that we said, "Not only is this cause kind of deserving of a hire in terms of how it scores in our ranking, but also we'd be able to..." Like, we think it makes sense to have a hire, 'cause some of the causes that are like too thin and too empty and too neglected, it's like we can't do it, like labor mobility. So the three causes where we've been interested in making a hire were criminal justice reform, factory farming and macroeconomic policy.

0:43:13 HK: Criminal justice reform, like that was the one we put first. The position is filled. And so now we're looking at the other two, and the factory farming we're prioritizing higher mostly because we just feel better about our odds of finding the right person there, than we do for macro. Like we've had a lot of conversations about what would a good person look like and what are our odds of finding them, and that's been a big factor for us.

0:43:35 HK: Factory farming, for those of you who don't know the cause, animals are treated really badly in factory farms. Those are the farms where people produce sort of like the meat that we eat. And there are a lot of animal welfare organizations, but they tend to focus on things like strays or endangered species.

0:43:54 HK: Arguably, if you put value on animal suffering, most of the animal suffering is actually going on in factory farms, and so, trying to do something about that... I think it's not terribly crowded. It's like a moderate, but pretty small field, and the importance is highly debatable. But I think we're kind of moving toward feeling that we're going to give a cause a high score on importance, if it scores high under a reasonable and defensible philosophical world view. We don't want to get stuck in one philosophical world view. We don't think that's consistent with the kind of organization we're trying to build.

0:44:30 HK: So I think factory farming, like some people would say animals shouldn't count at all, some people would say they should count equally to humans and this is the most important cause ever by a factor of 100. We're going to say neither. We're going to say, look, this cause is arguably very important and it's not very crowded and that makes it a good cause. Then I think it's reasonably tractable and there's some state and local opportunities, and I think it's not the world's most polarized issue in some ways, like there is... Polling tends to add people saying, both parties that animals shouldn't be treated terribly on factory farms.

0:45:01 HK: So yeah, that's another priority cause for us and one that we're looking to hire for at the moment. Again, any of these searches we're doing, I mean one of the things that we have been trying really hard to do is give ourselves options. So any job search we're doing, we can close the job search and say we're not hiring anyone, and that is just something for people to keep in mind. So we're running a job search 'cause we think we might be able to find someone great. If we don't find someone great, the position doesn't get filled, and I think that's the right way to be running things. Yeah?

0:45:28 S?: Is the scientific research also US-focused?

0:45:31 HK: No. So is the scientific research also US-focused? We have US Focus for policy for purely practical reasons, though I also think that the US is like, according to what I know, as good a country to work in as any. Because with policy, you're going for like small chances of the biggest wins you can get, and US, a win is really big. It's a big country. It affects the whole world. A lot of US policy issues are really global issues, like labor mobility. So that is a practical focus, I think quite a defensible one. Science, no, we don't care where the science is happening. We haven't found that distinction we need to draw. I do think that a very, very large chunk of the world's best science is going on in the US and it's often very useful when you're doing philanthropy to meet people and network with them directly. So there's often going to be a bias toward connecting with people, talking to people who are easy to talk to and connect

with. It's also true that like a lot of the world's best science actually happens in the Bay Area. So we don't have a US or Bay Area focus, but could I imagine a bias creeping in because it's like easier to meet someone when they're here? Yeah, and I think that's okay. I mean I think that's consistent with using our time in the best possible way. Yep?

0:46:45 S?: You mentioned from a global health perspective that GiveWell top charities are, as far as you know, very good. Do you recommend to Good Ventures and to other large scale philanthropists that they just fund those to saturation?

0:46:58 HK: Yeah.

0:46:58 S?: And if not, why not?

0:47:00 HK: So do we recommend to Good Ventures and other large philanthropists that they fund the top charities to saturation? We've gone back and forth on this a lot. Saturation would be a lot for an organization like GiveDirectly. Our general take has been for a very large philanthropist who's sitting on a big pile of money to say, it's good to give something. It's good to give something significant every year, and there's various benefits to that. I think...

0:47:28 HK: So I think one thing that's a little weird about top charities is I think there is a benefit to them getting a lot of money every year from a broad base of many donors. I think that their... Part of what we're trying to do with GiveWell top charities is change the way incentives work. Also, a charity is better off when it has more sources of support. So I think that's been one of the things that's made us hesitate to say just like totally saturate top charities. We have recommended big gifts to top charities.

0:47:58 HK: And so, our position at Good Ventures again I mean, you could... You could... This is another one of these debates. It's like you could easily end up on one side, and say, you should save all your money because the best giving opportunities are later or end up on the other side and say you should give every dollar you have today. But we've ended up saying just like when Good Ventures gives a big gift to our top charities, they do a lot of good, and it's good for incentives. But if they were to fully saturate, and push out all the individuals donors, I don't know about the marginal value of that. And so, I'd just as soon see that money saved for future giving opportunities. Yeah.

0:48:31 S?: So it sounds like you guys have done a lot of really thoughtful aspiration of what the unique role of philanthropy is in creating change. I'm curious as to how that's embedded in either the tractability or the crowdedness dimension of your considerations.

0:48:45 HK: Sure. Have we looked into the general, the comparative advantage of philanthropy, and how does that play into the way we look at causes. I mean, we've done a ton of thinking about what philanthropy is good at, and how philanthropy works. And now, it's some of the earlier stuff we did. So I've talked about how last year, we looked at a lot of causes. We did a lot of investigations. But before that was more stuff like reading about the history of philanthropy and getting to know some of the... Some of the basics. I mean, we did a lot of posts about policy and science that were just like basic stuff like, how does philanthropy affect politics?

0:49:21 HK: And yeah, I think that definitely affects our frame. I mean, when I think about... I do have a certain model of philanthropy in my head. I have a list of things that I've seen, reason to think philanthropy can fund successfully. I would broadly characterize philanthropy as having the ability to work on very long timeframes, take very big risks, do things that have to be justified with individual judgment, and with sort of judgment of one person or organization by another person because that's a thing that governments for example have more trouble doing. And I think that companies obviously have their own limitations. But I think working on long timeframes for the good of the whole world with high risk, and using a lot of personal judgment, I think those are advantages of philanthropy.

0:50:06 HK: And so, yeah, a lot of times when I look at how tractable is this? I mean, I think, well, so less tractability. When I look at how crowded a space is, I look partly... And who's everyone in the space? Who's everyone that thinks about this? But also I ask, "Is there other philanthropy in the space?" And if there isn't, are there opportunities to do special things by being unorthodox, by doing things differently, by just being a little more nimble, by taking longer bets. And that is how I think about it.

0:50:33 HK: Another thing that I've learned is that, I don't... I mean, I don't... I think philanthropy in some ways is less limited than a lot of people instinctively picture it. So if you decide that an issue is important, there are a lot of things you can fund to try and get the issue more attention. Get more people working on it. Make the world more ready when

interest in the issue goes up and when there's an opportunity to get something done. And so I do... I rarely look at an issue and say, "Yeah, this is a great issue. It's really important and no one talks about it. But there's nothing you can do as a philanthropist." I mean, I rarely say that. I think if no one works on a topic, if no one makes... Or not enough people make their career, thinking about a topic, and trying to do things, one of the things philanthropy can do is increase the number of people who work on something. A lot of... A lot of good philanthropy I think pays for salaries. And I think that's pretty consistent with the history of philanthropy. Yep.

0:51:30 S?: In your work on science do you focus mostly on biomedical fields or do you also consider physical sciences and engineering?

0:51:39 HK: Sure. In science, we mostly focus on biomedical or do we look at other stuff, too? We started with biomedical. This is kind of... A little bit of a pattern with us is that we often start in a universe that feels well established enough and mature enough that we can learn a lot. But within the universe, we pick causes. We like causes that are neglected and get no attention.

0:52:00 HK: So for biomedical, it's like there's a ton of philanthropy. There's a lot of people who learn from... There's a lot of stuff going on. It's already... People recognize it as an area that has humanitarian implications, and it's big, and there's a lot to think about. And if you were to change the way incentives work, you'd be changing a very large ecosystem. So that is where we started.

0:52:19 HK: We do have ambitions of looking at other fields. It's been just not... It hasn't been a high enough priority to be moving forward as fast as I wanted this year, so it is an official goal. This year, we do have an official goal of getting a very superficial quick glance at other fields and say, "Where else do we want to be working?" But I don't honestly know that we're going to hit it. But I did... The next field I looked at was computer science and specifically, machine learning and I had four or five hours of conversations on that and at some point we're going to put up the notes, and at some point we're going to talk about what we took away from that.

0:52:53 HK: So we do want to work on other stuff, but it's not the highest thing on the list. And certainly, I mean, I've had enough informal conversations with scientists who are kind of broad in thinking about these things. I mean, I don't... I have a decent sense of what we're missing, and I think we're missing some stuff. But it's not something that I haven't thought about at all, but it's something we haven't looked at as much as we'd like to. Yeah, Tom.

0:53:16 S?: So it sounds like most of the research processes these days focused on the experts and staff to support a particular project. I know in the past, there's been some talk about using the website as a forum for people to contribute ideas and contest potentially. And there's lot more work that's been done recently on open innovation tactics for learning, so contest and crowdsourcing ideas, and open source ideas. Now, I'm wondering if you guys are thinking anymore about integrating that into the process at all.

0:53:45 HK: Sure. So are... Are we thinking about integrating more sort of contests, open source ideas, decentralized search for ideas into our process? I'll answer the question a little bit more broadly. I mean, I think... I do have a rough set of things in my head that you can do when you care about a cause and you want to make some grants. And some of them are more time intensive than others. So, I think a lot of times the best easiest thing to do when you are interested in a cause, is you talk to the people who are already in the cause. And you say "We're really interested in this cause. We'd love to fund something." A lot of times that has results, and it has really quick results. And what happens is, someone who comes recommended by the three best people in the field has an idea, and the idea is similar to what you thought the field is missing, and then you fund it.

0:54:31 HK: This is efficient and this doesn't take a ton of time, and I think this is useful when we're looking across a lot of causes, trying to consider a lot of ideas. If you want to get more intense, there are other things you can do. So, I... We wrote a post called "Funder-initiated startups." Which is about when funders actually create organizations, or taking a leader role in creating organizations, and the history of this is a lot better than I would have guessed it would have been. And I did some reflecting on why that is.

0:54:58 HK: So, science is an example of where there is an organization that I wish existed. Right now, what I'm doing is telling everyone that and saying, "We'd love to fund this." And then if the right person comes to us and says, "I want to do this." That's pretty easy for us. The point of actually... More actively trying to create an organization. That would be a lot more work, and so, that's more of a measure that we would take if we were saying "We really need to make this happen. It's much better than other opportunities and it's better enough to take more time."

0:55:29 HK: Coming back to Tom's question about crowd sourcing de-centralized... I mean, we've looked at models like that and price models, contest models, and I would call that another high intensity intervention. I would call that something were you really... You really have to be very thoughtful about the structure. What is the deadline? What is it you're asking for? How is it that you're going to judge the entries, right? Because before people spend time in it, how are you going to... Because you want the best people to enter. And those are the kind of people who have the least time and want to make it's worth their time before they enter. And so, you need to be really clear about what you're looking for, how it's defined to the extent that you're going to be able to make sense of a lot of stuff that's coming to you without context. And so, I think that's something that... There is a time and a place for that kind of thing but I think it's not the first place we turn.

0:56:21 HK: You know, certainly the AI thing we are looking at is an RFP, which is a form of that, it's saying, "It's an open call. Please send us your ideas for working on AI safety." And I think that was the right call for that particular situation. But, yeah. I mean I think, look, all these things have a time and a place, but I certainly think that anything that sounds like crowd sourcing... I mean that sounds to me like a lot of our work to shape the crowd sourcing, and so, that just has to be taken into account. Yeah?

0:56:49 S?: Of the request for proposals for AI safety stuff?

0:56:52 HK: Yeah.

0:56:53 S?: How many of them in your eyes have a lot of merit? Are you generally impressed, generally unimpressed? Do you think there's a lot of room for funding there in the future here? Or...

0:57:01 HK: Sure. So, the question is kind of like how do we feel about the proposals for AI safety? I'm just not going to comment on that now, so that... I believe that the review of those proposals is currently wrapping up and then we're going to have a conversation, and then we are going to make a decision, and then we are going to write about it. So, I don't think this is the time to get into that. Yep?

0:57:21 S?: So, you guys seem like you spend a lot of time finding uncrowded, very impactful causes, but I'm sure in the course of your research you find unimpactful, very crowded cause. Is there anyone who is out there not trying to allocate new resources to the right places, but trying to deallocate resources from lesser impactful causes?

0:57:39 HK: Sure. Sure, so is there anyone out there, have we thought of getting money out of bad causes instead of money into good causes?

0:57:47 S?: Give poorly.

0:57:47 HK: I mean...

[laughter]

0:57:49 HK: Yeah, give poorly. Give less poorly. In general, that's not our model, that's not our focus, that's not what we do. And I think for good reason... I kind of mentioned what it looks like to have an influence on a philanthropist. You meet someone who's just starting to think they're open minded, and then you can help them spend their money in the best way possible, and that's a huge impact. To identify a cause as bad, takes just as much effort as identifying it as good.

0:58:15 HK: Arguably, a lot more effort because all the people you are relying on for information don't want to talk to you, or don't see you as an ally or don't think you're making good use of their time. So, it's more work and it's less good accomplished by doing it because then you say... Well, so then you need to find someone you can influence and anyone who's already giving to that cause is probably emotionally tied to it anyway. And then even if you did get them to switch, it's like that's still not as good as getting someone into a good cause. So, that's just not our model, but if someone wanted to try it, I would be interested to read their blog.

[laughter]

0:58:45 HK: So, it's 9:20. So, I just want to flag that and I think I'm just going to take one more question, then I'll hang

around for a while.

0:58:51 S?: Let me give you a quick comment. I've also thought you should send somebody to 18th and Valencia street, and tell all those people who are trying to collect signatures for bad causes who don't ever think about prioritizing anything. They just go to the GiveWell's website instead.

0:59:02 HK: I haven't gone down there and done that, but if you wanted to do it... I've had the thought certainly when I walked by, but...

0:59:12 Speaker 5: Only crowd sourcing people would do that.

0:59:14 HK: Yeah, crowd sourcing. Alright, so one more question. Unless there aren't any. Okay. I answered all... Yeah, go ahead.

0:59:21 S?: I was just curious about now that GiveWell or I guess Open Phil is moving into policy, just... It seems to me that people in general are more strictly sort of like... Values, ethics and not utilitarian when it comes to politics as opposed to private giving. So, have you run into any values conflicts with donors or former donors? Or internal conflicts where political disagreements have hindered GiveWell's policy work?

0:59:53 HK: Sure. Have we run into values disagreements where that's hindered our policy work? No. I mean not yet. I think we, obviously we've mostly been working with Good Ventures and I think we've largely seen eye to eye. I've also had a lot of conversations with donors and haven't felt a ton of tension. I think a lot of the policy issues we work on, it's like we've tried to... We've tried to look at issues where it seems clear enough to us, which policy change would be in the interest of humanity as a whole.

1:00:26 HK: And that is not generally what the political system optimizes for, so it's not that hard to find places where you can see what that change would look like and you can justify it. So I'm not going to say there's no disagreement, certainly there are people who have different political values and I'm friends with a lot of people who do have very different political values.

1:00:48 HK: But it hasn't been a major issue for us so far. And I think in general with Open Philanthropy, to put it more broadly, it is obviously a much more values laden project as a whole than GiveWell. Everything we do, it's a lot more judgment, guess work, risk. Kind of the whole point of it is to be bold and do what others won't. So we can't make it like the top charities where it's like everything we say is backed up and you can't really argue with it. Of course that's not actually how the top charities are, but it's further along that spectrum.

1:01:14 HK: And so, that is something we're cognizant of and that's why we want to separate the organizations. And our strategy is going to be to say "We are not objective, we are not neutral, we have values, we apply them as well as we can, we explain how we're applying them, we try to be open and transparent about how we're applying them, and then people who are excited by what we're doing can join in and people who aren't don't have to and people who are only excited by one part of it can do that." And that's our basic model and I think it's going to work. I think that will be enough people to have a big impact.

1:01:44 HK: So alright, well, thanks everyone and I'll stick around for awhile.

[applause]