

**GiveWell NYC Research Event, December 11, 2017 – 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.

0:00:09 Holden Karnofsky: Thanks for coming and thanks for being part of this event, the first Give Well session. I'm going to record and, Isabel's going to record just so [0:00:17] ____ that's just a reminder. And also a reminder if you have some questions, I'm going to repeat it for the recording. And if you don't want your question to be recorded or if you want it removed, just say so when you ask it or just email us afterwards and we're going to have no problem. So I'm Holden Karnofsky. I'm co-founder of Give Well and also The Open Philanthropy Project but as of this year, Give Well and Open Philanthropy are separate organizations. We still share an office. I'm still on the board of Give Well but otherwise, I don't work for Give Well. So I focus on The Open Philanthropy Project and I'm going to talk for a bit about just some of the basics, just to catch people up who don't know them. They're focusing more on what's recent for the Open Philanthropy Project and then I'll take questions just like in the first session. Basics, for those of you who aren't familiar, this will be some of you. Open Philanthropy started as part of Give Well, so Give Well initially was [0:01:13] ____ a proven, cost effective, scalable giving opportunities and we met Cari Tuna and Dustin Moskovitz a few years ago.

0:01:20 HK: And they were looking for how to give away a very large fortune in millions of dollars within their lifetimes. This, to us was a problem that in some ways is very similar to the one that Give Well tried to solve and in some ways, very different. The similarity is looking for how to do as much good as possible with your charitable giving and not really knowing where to start and not having a lot of intellectual content out there to help you do the most that you can. So that's the part that's familiar. And the part that is not so familiar is just the nature of the giving that Cari and Dustin were... The amount they were giving away and the time frame over which they were giving it and their ability to... Or their dedication to basically spending their lives on this. Cari was about to go full time when we met and she's now full time. So [0:02:06] ____ like that, I think has a slightly different situation from let's say, someone who is one of the original Give Well customers.

0:02:15 HK: So when we first started Give Well, it was based on me and Elie and what we needed when we were trying to give away money with a few hours to think about where to give it at the end of the year. And wanting something where we could really follow and understand the entire case. Working with Cari and Dustin we saw, maybe a little bit more of an opportunity to giving that could be in some ways higher risk, in some ways longer time horizon, and could require or involve doing things like building out their own staff, forming their own networks. Getting to the point where they trust certain people and have relationships with them, that allow them to do things that may be... To do giving opportunities that may be harder to explain and communicate to the outside world. And also just be able to do things like create new organizations rather than just funding those that exist, and to really transform existing organizations. That's where we're coming from, in terms of the overall approach of Open Phil and how we got to the causes we work on now.

0:03:12 HK: One of the key philosophies of Open Phil is this idea of [0:03:15] ____ space giving. The idea is, if you can take a lot... Well, you might look for giving opportunities where, for every 10 times you give you sort of fail completely and maybe even embarrassingly nine times, and then one does so well and it makes up for everything else. This is maybe a model that's familiar for venture capital and we think that this model is an interesting one for philanthropy that somewhat contrasts with the Give Well approach. Where Give Well tends to look for the things that are most dividends back. If you look at some of the biggest, historical philanthropic successes, some of them were so big. For example, the feminist philanthropy that really funded the research behind The Pill. The common oral contraceptive when the government wouldn't fund that, it wouldn't have funded that. Or the Rockefeller Foundation, funding research on improved crops that is widely believed to have led to the Green Revolution, making large countries kind of going out of poverty and becoming medium developed countries, and leading to probably over a billion people who didn't starve to death because of the Green Revolution or rather... Yeah, that's right.

0:04:26 HK: You can imagine that things like those are a part of philanthropy's history. You can also imagine that you could have a lot of failures. And if you have one success like that, you would really be compensated for it. So that's part of the general philosophy on Open Phil And in some ways, it's kind of a contrast to Give Well because I think Give Well looks for, a lot of times, the things that are most supported by evidence, things that can be understood, things that can be explained, things where you can trace the impact of your giving. And a lot of times, Open Phil is trying to just

swing so hard that the occasional win is good enough. This is a simplification and there are definitely points of overlap. Times when Give Well is funding things that...

[noise]

0:05:04 HK: And times when Open Phil is trying to stick with what's more proven or linear. But that is something of a theme of where you'll see the two diverge and take different approaches.

0:05:13 HK: In terms of how Open Phil has selected the areas that it works in, we spent the first few years really trying to build up a framework and choose areas to specialize in. Choose what we call focus areas. Basically the idea is, we believe one of the most important decisions a philanthropist makes is just what problems to work on. If you can choose the best problems to work on, then from there you can build your organization around developing expertise in those problems. And so we looked for problems or focus areas that we felt were some good combination of being important. Meaning, if we got a win, it would affect a lot of persons and it would affect a lot of neglected. Meaning we tend to look at who else is working in an area, and we don't want to work in an area that is overrun with other philanthropists.

0:06:00 HK: And then finally tractable, meaning we see some sort of path to impact. And if we can see causes that stand out on those criteria, those are the causes we want to enter. So we spent a couple of years just exploring all the causes we could work on, assessing them in importance, neglectedness, tractability. And then we spend probably about a year really focused just on hiring. So we knew what areas we wanted to work on, but we wanted to make sure we found the best possible people to specialize in those areas and to lead our work in those areas. And then from there, the Open Phil philosophy is really, you try to have the best causes, the best people and let those people lead the way from there. So there's not a lot more rules after that point. Our program officers usually are experts in their fields. They're usually among the most knowledgeable people about their fields and they become even more knowledgeable on the job, and we really let them lead the way. And we have what we call a 50-40-10 principle, which is that we want someone who, let's say, Chloe who works on criminal justice reform, or Louis who works on farm animal welfare. We ask that half of their portfolio by dollars be such that the decision makers, the people who approve the grants which is me and Cari, we can understand, explain and affirmatively agree with half of [0:07:18] ____, but only half.

0:07:19 HK: And so another 40% of that portfolio would be in what we sometimes call the okay category, which is like, "I can see how one might think this is a good grant, if one knew a lot more than I do." I don't personally totally buy it, but we're going to approve the grant anyway. And then the final 10% is generally discretionary. So a lot of times there's a very stripped down process and program officers are able to get grants through pretty fast without really convincing us of anything, as long as we don't see any red flags. And the principle there is that if you have the right cause, you pick the right cause and you pick the right people, from there we see a lot of value in actually minimizing the number of decision makers, minimizing the number of veto points and letting knowledge be distributed. Letting the leadership come from the person who knows the area the best. And we figure as long as half of our portfolio is something that we're really on board with, we can let the other half be stuff that we didn't have the time or expertise to fully understand the case for. But we've got the person who believes in it. And I do believe that a lot of the best giving in this kind of high risk genre is going to come from doing things that are hard to explain and they take a lot of context. And where you may need to really understand a lot of personalities and a lot of complexity. And so the more you can minimize that number of decision makers and veto points, the better.

0:08:38 HK: So that's the basic philosophy of Open Phil. Happy to answer more questions about what causes we work on and why. But for now, I'm just going to list some of the things we've been doing because it's now been a couple of years that we've been giving in the \$100-200 million a year range. We do expect that number to go up more over time, but for now, we're trying to learn a lot of lessons and get the best possible process in place and have a great understanding of what we're doing before we spend more. So some of the causes we work in and some of the things we've done. Farm animal welfare is an example of a cause that we selected and hired for based on one of the main themes of our giving which I refer to as radical empathy. Which is this idea that if there are certain parts of the population that today are not even considered to be people, and are not considered to have rights at all, and are highly marginalized, it's often the case, we feel, that our giving can do the most good by focusing on those parts of the population. And so if you look historically and you ask what are some of the best things you could have done, being early to various civil rights, feminist causes, I think before there was much agreement at all that those causes even were worthy of consideration at all, I think you would have done a really [0:09:58] ____ amount of good.

0:10:00 HK: And so a lot of the causes we work on today are trying to follow that same template. And some of them, such as criminal justice reform, focus on one marginalized person. Farm animal welfare is this idea that it's possible that

100 years from now, we'll look back at the way that animals were treated on factory farms and say, "Actually that was the worst thing going on at the time in terms of just how many creatures were being treated with incredible cruelty throughout their entire lives." And that may be something where we look back and say, "Actually, we should have considered those creatures be persons, to be things that could experience pain and that we should care about, and that therefore, that's one the worst things we've done." And so anyone who's trying to help them was doing one of the best things they could be doing. So that's the principle on which we believe that according to a certain set of assumptions, and I can get a little bit more into this later, how we treat very disparate world views and assumptions. But according to a certain set of assumptions, the importance of farm animal welfare is extremely high because you've got a lot of what we might want to consider persons worthy of moral consideration, just very, very, large numbers of them being treated incredibly cruelly.

0:11:14 HK: And then in terms of the neglectedness and tractability, we basically believe that we're by far the largest funder in this field now. Many fields we go into, we would never be the largest funder, not even close at this amount. But the amount we're putting in, we are by far the largest funder on farm animal welfare and we believe we've already seen a really large impact. When we came in, there were these corporate campaigns where people would try to embarrass various fast food companies or grocers into making what, in the end is a pretty economically cheap pledge to go cage-free. Doesn't raise their costs much, we believe it makes a huge difference for the chickens. And basically what people were finding is they were getting a little bit of momentum. Finding that, I think they had gotten McDonald's to pledge to go cage-free by basically putting this pressure on them for something that ultimately McDonald's didn't care very much about.

0:12:06 HK: We came in and we really tried to supercharge this and move it ahead, and funded all the groups to greatly expand their operations. Since then, since we came in, we've seen basically... Essentially every major grocer and fast food company in America make a cage-free pledge. So hopefully 10 years from now or so, any eggs you buy, they will all be cage-free and that will be the only way to get them. And we don't think we're the only ones responsible for that way. We think there was already momentum and all we did, we think was speed things up and provide support. But we think that the return was extremely significant in terms of just animal's health per dollar.

0:12:44 HK: It's kind of an odd thing, because I myself am not generally intuitively inclined to care about, let's say chickens as much as humans. But on the other hand, if you look at the Give Well top charities of, let's say a human death diverted for every few thousand dollars, and then you look at the farm animal welfare work, where we're estimating something more like 100 chickens per dollar that are now not going to be in cages. So it's enough orders of magnitude there to make you think twice. It will at least make you think, "Wow, it's possible under certain assumptions, and if we change our minds in certain ways in the future, or believe certain things in the future, that this really was outstanding work." And certainly in terms of just tangible impact, we believe we've seen it.

0:13:25 HK: Other things we've been doing in that cause is we've been trying to repeat this overseas. We came in a little bit on the tail end of the US cage-free stuff and that there was already a little bit of momentum, although it was definitely exponential to when we came in. But we think that in a country like China, that's an area where meat consumption by default, it's likely to rise a great deal over time. And getting that start, getting that foothold today, could make a wave of corporate campaigns like this or something comparable possible in the future. So we've been taking a much more long-term view overseas. But in Europe and Latin America and India and China, we've been trying to bolster the farm animal welfare movement and in many cases, put something there where there was nothing. We are the only farm animal welfare funder I think, essentially in China at all. So that's the other stuff we've been doing. And we've also been funding just research to better understand what kinds of things these corporate campaigns should be pushing for. So how could you reorganize some of these agricultural industries so that the animals are suffering less. And then you can build those into these asks around these future corporate campaigns.

0:14:34 HK: So that's examples of some work we've done. Criminal justice reform is another one of our focus areas. I mentioned that a little bit ago. We picked that one because out of all the US policy issues we looked at, this one seemed to really stand out in terms of [0:14:48] _____. We thought we could maybe see a path to impact. There's a bunch of reasons for that. One is that it's a state and local issue and so it's a place where you can pick your battles on this issue, unlike many other important US policy issues. And it's also just an issue where I think the winds are changing a little bit; state budgets are tight, crime rates are down, and there is both liberal and conservative support in many cases for cutting incarceration, and we do believe there's too much incarceration in the US.

0:15:18 HK: So that's another area where there would be a major focus, and we believe we've also seen some impact there. We supported the Alliance for Safety and Justice, helped them expand faster in Illinois, and there was a bipartisan

reform bill that we believe is going to lead to several thousand people a year less every year in prison. Which we think is incredibly inhumane and a big impact. We've also supported the Closed Records campaign in New York, that a few months after it started there was an announcement from the mayor that New York City does intend to close records. The plan has been outlined so far, I think it's not as aggressive as what we were hoping to see. But those are some examples there.

0:15:58 HK: Other causes we work on, we support scientific research funding. That is something that in past events I've often kind of said, "We're still figuring out the framework for that." Now I think we have a pretty decent framework. We have a couple of senior science advisors. We have a list of areas of interest, and we generally ask them to be opportunistic and look across all of the things they are seeing to find especially tractable, especially exciting scientific opportunities on some of the things we've done. We've made a major grant for gene drives, research on gene drives which we think could be one of the most likely paths to fully eliminating or eradicating malaria. And we also tried something this year, where basically the National Institutes of Health has a special program for high risk, high reward scientific research, and it's very small. They fund almost none of the applications. We essentially managed to get almost their entire reject pile, just all the things they didn't fund, and we went back through it and looked for things that we felt would be interesting to fund. And we did fund four projects. One of them for example is a shot at a cancer vaccine that I think is at a level of high risk, high reward that I think is even too much for Pharma. Very hard for putting up the best in one place and making a profit on that, but we thought it was worth a shot to fund.

0:17:16 HK: And then the final general category of what we work on in addition to trying to focus on marginalized populations, we're also very interested in civilizational or global catastrophic risks, and potentially other [0:17:32] _____ factors too. So things that could be upside possibilities as well as risks. But the basic idea is we believe that the world today is a little bit more interdependent and interconnected than it's ever been before. Probably be easier for humans to drive ourselves extinct today than any point in the past and it's only getting easier over time. And so one of the angles that we take to try to do a lot of good per dollar is, if you could make even a relatively small contribution to prevent a global catastrophe that could result in a new equilibrium, such as humanity going extinct or such as some very bad and very durable and irreversible outcome, then you're not only helping the people who are directly affected today, you could be helping many, many, many future generations.

0:18:20 HK: So that's another spin, that's another take on how to get a lot for your money in some sense. Major causes there. We have not seen... These are longer term, lower probability plays, but we do focus especially on bio security and [0:18:34] _____, and on potential risk for advanced AI. Those are basically two things that are at the top of my list, if something was going to be catastrophic in a really global, far-reaching, irreversible way. Probably come from one of those two areas. And a lot of what we've done is try to take what we consider the biggest, the worst cases, put them on the map, get them taken seriously, and build up fields of people who could provide expertise on them. So I'm happy to elaborate on all that in response to questions.

0:19:03 HK: Before we do that, I'm just going to talk a little bit more about just what we're trying to do organization-wide right now. So far we've chosen causes, we've hired people, we've made some grants, we've set up our basic processes around making grants. But some things that I think are really on the IMAP for going forward, and some of the things we're trying to do across the organization, one mode is just continue to heavily critique ourselves. So a lot of these causes that we're in rely on assumptions of one kind or another or beliefs of one kind or another. And the deeper we get into a cause, the more we commit to it, the more money we put in. And in general, I believe that in this kind of philanthropy you're much better off with a 20-year massive financial commitment that is very, very intense and allow the whole field to rely on you, than you are going piecemeal or making a new decision every year. And so the deeper we get into these causes, the more important it is to question our assumptions and beliefs.

0:19:57 HK: So for example with criminal justice reform, we went into the cause with a view that the US has just really outlier levels of incarceration and that it's probably too much. That's conventional wisdom among the most informed people we talked to. But David Rubin, who you may have also heard at our last session did a really thorough investigation of the literature on... Let's say we did reduce incarceration, should we expect that to cause problems in terms of increased crime? And so he did a really, just one of the most, I guess probably just the most ethic literature I've ever seen on anything.

[chuckle]

0:20:34 HK: Basically trying to find high quality studies, "Does incarceration affect crime? And what direction and what margin?" And he took all the studies that he could and found the data for them, reran them, reconstructed

everything they found, often found weaknesses. I feel like just about half the studies he looks at in this way completely fall apart. And try to synthesize the conclusions and did come out feeling that there's not a ton of, at least in the current literature, not a ton of argument for the idea that incarceration at the current margins is doing really anyone much good. He thinks that, basically his best guess is that the net effect on crime of reduced incarceration at this time would be zero because prison has some criminogenic effects as well as some crime preventing effects that, he thinks that at the present margin, those are probably coming somewhere close to netting out. And then we've done similar things in other causes, so we're doing an extensive [0:21:33] ____ threat level for pandemics, trying to figure out what the biggest threats are and how big they are. And on the timelines, doing some work on just trying to think about what can be known about advances in AI in the future, and what kind of capabilities one might expect.

0:21:50 HK: And then we also did a report on, "Is the cage-free reform really better for chickens?" Someone challenged us on that, and we wanted to look into it, and determine just how true it was. So those were some of the things we're doing. And then other things we're working on, I think we want to put more work and time into getting really good at self-assessment and looking backward, analyzing our impact, looking at it critically. So when we think we got a really good win, we kind of analyze it critically with the stories, the same way that we've done our history of philanthropy case studies, and see if it's really what we thought it was. We want to start booting up that process. And we also wanted to get better at comparing what our progress is, and what impact we've seen to what we predicted at the time of making the grant. And one thing that's helped us with that is that for a very long time now, every time someone suggests a grant, they need to make predictions with probabilities and due dates on them.

0:22:43 HK: And so then we can go back when the grant is due for renewal and say, "Which of these predictions have come true or false by now?" And that does help frame things. And then finally, a major project for us right now is cost prioritization. So that's figuring out how much money we ultimately want to put into these many very different goals with very different metrics, and that can be a little bit of [0:23:02] _____

[noise]

0:23:03 HK: Challenge. Because at first glance you would just say, "Why don't we just fund all of the stuff that does the most good for dollar and stop funding when we run out of money?" But your estimates of how much good you're doing per dollar can be very sensitive to things like whether you want to count helping chickens or not, or what fraction you want to count it, as compared to humans?

0:23:23 HK: And also very sensitive to things like, "How many future generations are you helping if you do help prevent global catastrophic risk?" Sometimes when the count is that sensitive and the question is that hard to answer, we think we need to do something other than just running the numbers. We think a lot of times we need to just say, "This is a really tough fork in the road, and we're going to put some resources in each branch." And figuring out how to do that in a principled way, I think is something that we're struggling with right now. So that is who we are, what we've done so far, and I'll just spend the rest of the time taking questions. Yep.

0:23:56 Speaker 2: Do you approach allocation the way a financial manager would? Where you have... I think you just said about the [0:24:03] ____ in the beginning, do you redistribute based on... Okay this seems like it's using a lot of money [0:24:11] _____. This one seems like it's making way more than we expected, and it was only supposed to be at 10%, so we're going to reallocate. How do you...

0:24:20 HK: Sure. So do we handle allocation like a financial manager would? Do we look at performance and reallocate money based on that? But most of the answer to this is we don't know yet. So most of the answer to this is, two years ago we were just starting to get in scale, and we had just starter budgets. So we thought we'd be big enough to see how things go, see what kind of impact we were having. We're only now even starting to be capable of asking how our different portfolios are performing relative to each other, and still not really. I think two years is just, in my opinion it's basic philanthropy. Two years is just not long enough to know if you're having an impact yet, 90% of the time.

0:24:56 HK: There are a couple causes where we think we've had some impact. There's one cause, farm animal welfare, where we think we've had a lot of impact. But I don't think it's really time yet to start judging these portfolios in a definitive way. I think we're figuring out that framework, but I do think, ultimately, that we want a cause allocation that pays attention to how much good we think we can accomplish in each cause. And in theory, what we want to do is do the most good, total. Which means that we like to take money out of causes where we're doing less good per dollar, put it into more causes where we do more good per dollar, and of course, there are some of those issues that I brought up before, where some of the numbers can be very sensitive to these questions.

0:25:33 HK: I think we're still feeling our way along with that. One thing I will say, is that an important difference that I think can be... It's easy to miss, in philanthropy versus for-profit investing, is the idea that I think risk aversion makes a lot less sense in philanthropy than in for-profit investing, and that can actually change a lot of things. If you're doing for-profit investing, someone will tell you to diversify. Which means, instead of just spending all of your money on one stock, you should bet it on a bunch of stocks to lower your worst case scenario.

0:26:05 HK: But actually, if you weren't worried about going broke, if you're doing this sort of for the world instead of yourself, and you care about the total amount of wealth in the world, and you weren't worried about going broke, you might actually be better off just putting all of your money into the one stock that you think is the best bet. It would be scary from a personal point of view, but from a total amount of wealth in the world view, it might be better. I think in philanthropy it's just odd, and I think it's easy to miss, but there's fewer reasons to diversify, and I think there's more of a default position of, if the best thing can take all money, then that's what we should do, period. And then I think there are actually some considerations in favor of diversification, but not as many, and it's not as obvious. And that's one of the things that we're really debating internally right now, as we try and set the budgets. Yep?

0:26:57 Speaker 3: So does all of your money come from that one donor? [0:27:01] ____.

0:27:05 HK: Sure. Does all of our money come from one donor? Cari Tuna and Dustin Moskovitz are the funding partner that we're focused on right now. Dustin is the co-founder of Facebook. And at this time the level of giving we're at is lower than the level of giving we eventually want to be at. And so for that reason, we're just not very focused on going out and raising money, finding other funding partners. There have been... I think it would not be quite right to say they are the only funders. There's definitely cases where we.

[noise]

0:27:36 HK: Money one way or another from another donor. A lot of times it's more a matter of saying, "Hey, here is a cause area." Or, "Here is an organization that we recommended." And someone will give to it. We actually have a blog post coming out about that, I would say this month that's going to talk in a little bit more detail about just what we've seen from external partnerships. So I would actually say that there have been cases where, in one way or another we're talking to another donor or making recommendations, they're following the recommendations, and that is money moved. But in terms of process wise, the process is mostly just go through Cari and Dustin for now, because we have the longest running relationship with them. Cari's in the office. We are able to have a process with them in the loop. And then we are not really in need of additional funds at the organizational level right now, because of the fact that they're not as high as their giving eventually needs to be. That said, there's some causes where we've given someone a budget, and then that person has been much more aggressive in raising money from external donors just to supplement what they're getting from Cari and Dustin.

0:28:43 S3: Cool. So if you're... What's... I guess your time horizon is to spend it before they die, or?

0:28:50 HK: Yeah, I mean that...

0:28:51 S3: What does that look like in terms of a break up of spending money now, versus...

0:28:56 HK: Yeah.

0:28:57 S3: A few decades from now.

0:29:00 HK: Sure, so what's the time horizon for spending the money? Yeah. Cari and Dustin want to give away the vast majority of their wealth within their lifetimes. That does imply certain things about how high giving will have to be at its peak level in order for that to happen. And definitely, the current number is lower than it will be at peak. I don't think we're committed to any particular schedule. So a lot of what's going on right now, is I'm thinking we are still learning, and I think this would be the wrong time to go up to peak giving. I think it's very hard... It's much easier to grow a budget than shrink a budget in philanthropy because you've got fields and organizations relying on you. So I think we want to put in more thought and have more of a sense that we nailed every aspect of our work before we aim for the peak. That said, I think once we feel we really know what we're doing, I mean there won't be any rule against spending it all in one year, if we think that's the best thing to do. My guess is that in many cases we want to reserve some capital: A, to be able to meet long term commitments, stay in the field for a long time, be reliable.

0:30:00 HK: I think a lot of what [0:30:01] ____ are able to do is help fields grow. To help the field grow, people need to believe that there's a career in the field, they can work in it, they get constant support. So I think there's a lot of good reasons to make sure you have enough to stay around, for a while, not for forever, and then I also think there's just generic, "It's good to keep some dry powder in case something outstanding comes up". So I don't know exactly how it's going to work, but my guess is we're going to spend a little longer making sure we really know what we're doing, and then we'll have more of an idea, and then probably bring the level of giving up significantly, and then slowly down. Yeah?

0:30:38 Speaker 4: [0:30:39] ____ your current success with impulse funding on farm animal welfare, but there were orders of magnitude more wild animals than farm animals. Some of them may be conscious. Does corporate developments to funding in that area... I know there's not much to fund at this point, but maybe someone could get the ball rolling and in the future there may be more things.

0:31:05 HK: Sure, so farm animal welfare versus wild animal welfare. There's more animals in the wild than on farms. Have we thought about trying to improve the welfare of wild animals? And certainly, it does seem like there are certain cases in which, whether for human behavior reasons or otherwise, definitely there can be cases where you would look at wild animal situations and say, "Yes, well it could be improved a lot." So are we interested in that? Like you said, I think we do not receive a lot of funding opportunity right there right now. I think there is actually a huge difference between wild animal and farm animal welfare in terms of just what can you do. Farm animal welfare is a case where incredibly abusive things are happening and it's really... To anyone who hasn't read a book about this, or watched a video about this. The objection here goes way beyond the idea of killing animals for me, it's how they're treated throughout their lives, in this incredibly inhumane way, that just it really just looks like it's someone trying to squeeze every penny of cost out of it, and not give any regard at all for the comfort of the animal.

0:32:11 HK: And you make people aware of that, just the cage free stuff a lot of it's just you kind of highlight to ask the company's customers that the chickens are in tiny cages. People just don't feel good about that, and then there's some embarrassment, and there's a change in policy, a lot of animals are affected at once. So there's this opportunity sitting there that this behavior is going on, and really people I think should not be proud of, and most people would not feel comfortable about if they saw it and knew about it. And it's a human run system, and so you can get changes made at scale. I don't see anything like that for wild animal welfare, and I haven't really seen anything to fund that looks comparably effective.

0:32:50 HK: If some giving opportunities came up, I think we would just look at them, and think about, what is this? Is this a good... Would this help a lot of animals for a little money? Does this have other issues? I think sometimes with wild animal welfare you would worry about interfering with natural habitat, we'd have to weigh that against whatever the opportunity was. So there's certain categories of things where maybe cutting back on human abuses that lower harms to wild animals, there maybe other things that are intrusive and you have to weigh that against the benefit you're seeing. At this point, I mostly don't [0:33:23] ____ about it, it's the kind of thing where we're aware of it at a low level, and if there were more shovel-ready things to do, places where it checks, I think we add that to the conversation at that point. Do you have a follow up?

0:33:36 S4: Yeah, so, I agree there's not much to do in terms of practical change right now, but [0:33:44] ____ developmental on research and it seems like they don't feel wild animal welfare [0:33:51] ____.

0:33:55 HK: Sure, so one of the things you could fund is research on wild animal welfare, and just people thinking about is there may be stuff we could do. And that is something where, again, actually even at that little level I'm not aware of what I consider opportunities that I looked at and said, "Boy, that looks like a good return and addresses all the issues one might have." But again, if it came up, we would assess at that point.

0:34:21 Speaker 5: Kirk, thank you so much. My question is also with the animal welfare area. A recent grant with The Guardian to fund journalism. So it's about animal welfare, but it's also about the role of journalism, what folks see as the role of journalism. And have you funded other journalistic projects outside of the animal welfare area and just how do you think about that and also how do you think about measuring the impact around journalism?

0:34:45 HK: Sure, so what do we think is the role of journalism? How do we measure impact? And this is in reference to a recent grant being announced to The Guardian to fund... So I'll try and explain the grant. Basically, The Guardian is a mainstream media publication, they've got a significant interest in farm animal welfare, but journalism, I think is an

increasingly difficult enterprise financially. And these are not the articles that are the easiest to support in the typical financial ways through perhaps corporate sponsors, through advertising. And so what we try do is provide support that enables them to cover farm animal welfare a lot more than they would be able to otherwise, to hire an investigative journalist. And to go do some of the work that we think is one great example of how journalism could shine a light on horrible practices, but doesn't seem as financially viable right now as a lot of other things journalism can do that we think are less impactful. So yeah, we are trying to support that. It is an experiment, and so we have not done anything else like that, The Guardian is the only thing going on right now and we do see it as, "Let's see how this goes." Certainly, The Guardian is planning to share with us whatever data they have on... Well first up, we're going to look at how the articles come out, and are the articles good? And are they shining a light on important information.

0:36:01 HK: And then we're going to look at are these articles getting attention? And there are some measures of that, shares and likes and all that stuff. And Guardian's going to share that info with us. So that will be some of how we evaluate it. We don't have any one formula for evaluating a grant. A lot of things we do are just kind of... It'll be hard to observe the impact within the next five to 10 years or maybe ever. But we can at least get an intermediate sense of, is quality journalism happening? Is there a buzz about it from communities of people who seem to be in a place to do something about it? And then there's the more generic measures of traffic. So we're going to see how it goes. And I could imagine doing similar things in other causes but we aren't at the moment looking at anything like that. Our view is just that it's a good combination of an issue that we thought journalism could do a lot for and also they were really excited to do it. And they wanted to hire other people. It wasn't something where we were trying to push it on them. It was more, we were trying to make it possible for them to cover that moving on.

0:37:02 S?: So in addition to the different areas that you describe, Open Phil also gives grants to Give Well's top charities?

0:37:12 HK: Yep.

0:37:13 S?: So I'm wondering how does the amount of money that Open Phil has compare to the amount of money needed to say eradicate malaria or things like that, solve some of the problems that Give Well's top charities do? And is Open Phil intentionally holding back from that because there might be better opportunities? Is it because there's not enough room for more funding right now?

0:37:41 HK: Yeah. So what is... Open Phil does recommend to Cari and Dustin that they make grants to Give Well top charities. And that has been an evolving thing over time. Because originally when we met Cari and Dustin we just were Give Well and that was the first thing we were doing and they were excited about that they wanted to give to. And so it's not as though I sort of came to Cari and Dustin saying, "Hey, there's this thing called..." That's not a history, but the part that Open Phil plays in this is we do try to help Cari and Dustin figure out how much each year to give to Give Well's top charities relative to all the other things they are compensating. And this is something where our thinking has had a lot of ups and downs and evolved a lot. One thing to just keep in mind is the roof for the funding on the Give Well side has been going up over time. So now it's in the several hundred million dollar a year range.

0:38:29 HK: And that we're... Open Phil giving nine to the Give Well top charities. Open Phil giving is like a \$100 million or \$200 million a year. So you can fill that gap. Well at least for starters it would be a doubling or more on improbably bad numbers that build more over time. So that's something to know. And I think one of the principles we followed here is just kind of, throughout we followed a principle of some degree of caution and not wanting to take irreversible actions. And so there was a time in previous years when we had a very tentative guess that we actually thought Give Well top charities were probably better than alternative uses of that money in the sense that, whenever we're contemplating a grant we kind of compare it to the last offer. Because you can make this grant or you can save the money. And if you save the money then you'll sort of spend it at the end after everything else that you funded. We used to have a view that the Give Well top charities were probably better than the alternate use of that money and yet we still didn't recommend fully funding them because we were very unsure about that.

0:39:26 HK: And what we didn't want to do was start to create a precedent where every time Open Phil or Cari and Dustin like something, they fully fund it and everyone else should stop funding it. So I think it wasn't a good place in terms of raising and growing an amount of money from individual donors. And I think if we had recommended to fill that whole gap, I think the individual donors quite rightly would have stopped giving there and I think it would have been hard to resurrect that. It would have been hard to resurrect that whole situation. So that was the situation at a point in time. At this point in time the situation is different. So at this point in time we basically feel that according to some world views and sets of assumptions about the world, the last dollar is actually going to do more good than Give Well's

top charities. And examples of those world views I think if you're a person who believes, for example that you should care about chickens 10% as much as human for example.

0:40:27 HK: They probably think we can spend that last dollar on something better than Give Well's top charities. And probably on animal focused things. Similarly, if you're a person who has a particular aggressive set of views about future generations and the importance of future generations and the number and the nature of future generations, then you might think that the last dollar spent on global catastrophic risk prevention, in the end is going to be helping more people per dollar expectation than Give Well's top charities. And if you reject both of those, there might still be other views that we've had out there. There is some set of views that we think is reasonable that says actually Give Well top charities are better than the last dollar. And so what do you do with that, and then it becomes a matter of just saying, it's sort of the world view application problem, which is how much money you want to put down this fork in the road, this path in the fork versus that path in the fork.

0:41:17 HK: And so the concept is we want to divide up the money and some of the money is going to follow one world view, one set of assumptions. Some of the money is going to follow another world view and according to our best understanding and our best estimates of how that allocation's going to end up when we're done with it. That is how we determine the Give Well allocation this year. So actually, this year the number does not... It's not based on the funding gap. It never has been based on individual donors. We've never done the thing where we say, "Let's give less because the individual donors are giving more." Never done matching either. In the past it's been a function of funding gap and this year it's not. This year it's just a function of how much money do we eventually expect to spend under that set of assumptions that favors Give Well's top charities. And then based on that we worked with Give Well to determine how much should come out this year based on how much is going to come out eventually. And we also did some rounding, so the number ended up being \$75 million, which is not nearly enough to close the gap, not even half, so that's the situation.

0:42:22 S?: You mentioned some follow up research into whether going cage-free was effective. How has that been going?

0:42:30 HK: Sure. So we followed research on whether cage-free is actually good for chickens. We put a blog post up about this and it goes to a write up, so there's a lot of detail about this on the web. We concluded mostly that our original view was essentially right. We're less confident in it than we were at the beginning so I think we did learn some things that made us pause a bit. I think there's a little bit of ambiguity, there's a possibility that mortality goes up in the cage-free systems, and there's a view that that's mostly because the cage-free systems... It's like when you transition to them you have to learn a whole new set of practices, and so it could be that things get worse for a period of time and then get a bit better, and it could be that the mortality's even higher in the long run because the chickens are kind of more free to interact with each other. But we think that's probably offset by the chickens getting to do a lot more of the sort of natural animal behaviors.

0:43:22 HK: And so overall we ended up feeling like the state of the literature is not what we wish it was. The thing is not totally unambiguous, but we do think that the cage-free reforms are, in the end, roughly as good for chickens as we thought they were. I do want to clarify, we don't think cage-free systems as used on industrial farms, we don't think those are humane; we don't think those are okay, we don't think those are reasonable. But we have been pushing for things to get better one step at a time, and that is not going to be the last act as far as later events, but also for now we're also moving on to some of the corporate campaigns to other asks, such as broiler chickens and other [0:43:56] ____.

0:44:00 S?: When you think broadly about cause allocation, how do you make trade offs between length of suffering, severity of suffering and mortality?

0:44:06 HK: Sure. So when you think about cause allocation, how do we make trade offs between length of suffering, severity of suffering and mortality? There I think, Give Well has probably done more work than Open Phil on those particular trade offs, and so when needed we're able to borrow on their work where a lot of it is just laying out the tough questions you have to answer, having different people who have been paying a lot attention, so people on staff, putting their numbers, their assumptions, taking some aggregates, statistics like medians. And those people who are filling the numbers are also trying to work off research showing how the general population might value those numbers, so we're able to borrow that when needed. It hasn't been as big... That particular set of questions has not been as big an issue for Open Phil. More of the kind of deep judgements have to do with the number of future generations and the relevant importance of animals; just don't think those trade-offs have been really crucial to the decisions we've made so far.

0:45:04 S?: Yeah, so the last couple causes that you mentioned AI risk, global catastrophic risk, something about them just doesn't sit right with me and I'm just going to think out loud for a second 'cause I don't know exactly what it is. But they're different than a lot of the other causes you work on in that they're not neglected, or they're not as neglected in that the incentives are aligned. It's not like the people who are doing the damage wouldn't suffer in the same way. Whereas in a lot of other causes the suffering is due to people being disconnected and the incentives aren't aligned, basically. Whereas in this cause everybody would kind of suffer equally and everybody's like in the same boat. And just something about them just doesn't sound like philanthropy to me, and I don't know exactly what that is but I was hoping you could talk more about that?

0:46:04 HK: Sure. So the goal of catastrophic risk causes is if something doesn't sit right, they don't feel philanthropic, and I do think a lot of them... I have some similar intuition to what you're saying, not saying it's the same, but I think some of it may come from the observation that global catastrophic risks are not about helping people who can't help themselves, it's a lot about working on issues that affect everyone, and that raises the question "Well why are these neglected?", or "Are these neglected?" And so if you have something that might wipe out civilization, it seems like something everyone should care about, and so why is it neglected? Is it neglected? Why are we working on it? , and can that really be called philanthropy if it's something that sort of benefits everybody equally and isn't about trying to benefit especially marginalized folks. And so I do hear all those properties being true.

0:46:56 HK: That said, I think global catastrophic risks are extraordinarily neglected; I think there's really solid arguments that they're more neglected than maybe some of the other work we do, so I think weirdly I feel quite confident that the potential risks for advanced AI we're worried about are far more neglected than criminal justice reform, even though I think criminal justice reform is pretty neglected in the scheme of things. So why is that? One thing is something about climate change as just sort of an example, or a point along this spectrum, and we've done some climate change work as well. In theory climate change affects everybody and everybody should care about it. In practice, a lot of the work on climate change mitigation is coming from non-profits, it's coming from altruistic folks, and I think you could try to explain why that is. I think some of it is just that it affects everyone, but in a very diffuse way. And so if I were running the US Department of Defense, probably I'm more worried about various things that might embarrass the US, or just hurt it in a more proximate, higher probability way.

0:48:04 HK: And in some sense, all of civilization going under is not a lot worse, from the perspective of the military doing it's job than, let's say, a problematic military loss or something like that. And I think, even though everyone has some incentive to care about global catastrophic risk, the amount that any one person or any one institution or even any one government is really threatened by this thing, it could be a very long way off and a very low probability, is actually quite minimal. And a lot of times higher likelihood but less large problems can occupy their attention. So I think that is certainly the case for climate change. I think it's even more the case for worst case pandemics and even more the case for potential risks from advanced AI. And yeah I think I'd feel a little bit odd if all of the philanthropy was going to this kind of thing for reasons you say. But ultimately I look at the situation and then I think for whatever reason, especially why, but for whatever reason, really these are very neglected. Really they are not getting much attention. And really, we believe we make a huge difference in the size of the field of people who pays attention to these things and if our goal is to help the most people for the least money that actually does seem quite competitive with everything else we're doing. Yep?

0:49:22 S?: Do you know of the investment in for-profit corporations? And if so, which one and how much? And if not, maybe something that you would consider?

0:49:32 HK: Sure, so the question's about for-profit investing. Do we do that? Would we do that? So in general, as a general matter, we think of our expertise, our specialization as finding ways to help people, as opposed to making profit for ourselves. And so any time the main case for something is make a profit, we think that there would other people in better positions to evaluate that than we are.

0:49:57 S?: There could be for-profit corporations that would make the world a better place for everybody and are neglected for some reasons that would not get...

0:50:07 HK: Sure. So there could be neglected for profit corporations. And I think as a general matter, I think that there's a lot more people interested in for-profit investing than in non profit. And so I think if we're trying to fund the things no one else will fund, or do better than other people, I think we have a much tougher field to contend with in for-profit than non profit. So generally, we focus on a part of investment that is not paying for itself. That said, there's no reason we can't do for-profit investments and we have, but just when we do them, we would tend to say, "We think that

this thing is more good for the world than it is profitable." Which means we're willing to sacrifice some financial return on it. Which means that we think of it as a generic investment plus a grant. That's one way to think of it. So if you invest in a company that you think has, let's say, a 10% chance of getting you a five X return, you might think of that as 50% just normal investment or something, and then 50% grant.

0:51:11 HK: And that's how we would tend to think of it. So we don't really try to have an edge on assessing the financial side but when we think the benefits of society are greater than the profits, we may make an investment and call a grant that size. And so an example of this, we do have an investment in Impossible Foods and they're a company that's trying to make plant based burgers and other things that are really convincing and that really taste like burgers and that could possibly replace burgers and... Here's a case where I think the amount of good to be done, is in some sense greater than the amount of profit to be made. You certainly, if you could pull that off you would make money. You would sell hamburgers but you wouldn't be getting paid for all of the animal suffering you're averting as well as some of the climate change implications of this and other implications. We have done it, but in general... It's a kind of thing where we have our causes, we look for the best things to do, if the best thing to do happens to be an investment instead of a grant, we'll make the investment. But we don't look specifically for investments. Yeah?

0:52:11 S?: Are you interested in commentary or anything about individual giving and the trade off of sure things that will give me at least some sense, in the Give Well sense that I've done some good versus the high risk, possibly high yield giving that you're able to do. What kind of advice would you give to individual, let's say, regular folk, who are trying to think about how to do some good with their giving relative to those two models?

0:52:43 HK: Sure, what advice do you give to individual donors and talk more about the relationship between individual donors and the mega-donors, like Cari and Dustin. I do think it's easy to oversimplify this stuff, so I do want to be clear. I think Give Well does do some work that ' quite high risk, long term, high reward and they do serve some very large donors, including, Cari and Dustin and I think that makes sense. But I do think of the two organizations as sort of occupying different places on the spectrum. Where in general, what Give Well recommends is going to be something where the full case is going to be written up online. You can check every component of it and you can have a certain sense that this is all holding together and making sense. And I think in general, that is going to be an especially attractive thing for people who are trying to make a giving decision relatively quickly. Whereas, I think Open Phil tends to occupy the niche more with some of the wackiest, the most ambitious, the highest risk things, and a lot of times these things are the most appropriate for people who really, really trust the people they're dealing with. Often in a way where they may want to build that trust up over the course of years. So it's not a totally thick line. And to illustrate that, price of the individual donors is, I would kind of...

0:54:00 HK: Did we end up running... Do we have a this coast of three kinds of individual donors or is that not up?

0:54:06 S?: No but there's, are Give Well's top charities the best option for every donor blog post, which I would highly recommend, since I think it exactly addresses this question.

0:54:13 HK: Great, so that blog post... I think the possibilities of online, giving to Give Well top charities is one great option. I think if you are very insistent on doing things that are in the more high-risk bucket, then I think you have to ask yourself, is there someone you trust to do some things that are going to be harder to explain and less checkable? Well, and then you can fund that person or that organization if so. And if you are looking to fund the sorts of things that Open Phil does, there are these funds called the effective Ultras funds that were set up by another organization but a lot of them are administered by Open Phil's staff. For example, there is an animal welfare fund and that is a fund that Louis, who runs our farm animal welfare work, he will draw from that fund to fund things to supplement his Open Phil budget. And so that option does exist but you're not going to get the same... There's also our future fund that has the future generations idea, also with an Open Phil staffer, but you can give to those if you want. You're not going to get the same level of transparency and sort of the whole case be online that you get from Give Well, the level of trustees. I'll take a couple more.

0:55:32 S?: What organization does the effective altruism funds [0:55:34] ____?

0:55:35 HK: Center For Effective Altruism. I'll take two more.

0:55:39 S?: Sorry, you spoke about the differences between worldviews of one who would be more inclined towards Give Well's funds, the causes of these [0:55:47] ____ worldview of someone who would be more inclined toward Open Philanthropy, Open Phil. I noticed that there's one particular person who is involved in founding both of those

organizations. How much of those two different worldviews reflect the image of your own worldview? It seems interesting that they seem to be going in two different directions and that you were involved in both of them.

0:56:06 HK: Sure, so my own view on the two going different directions... Quick point of clarification but there's a lot of people who've been heavily involved in both Give Well and Open Philanthropy. So in terms of answering your question, yeah, I actually like both. I, generally, in this space feel that extremes are good and the middle is my least favorite place to be, so in my view of... If you want to do something where you're being very self-skeptical and you're making sure everything checks out, I think Give Well does a better job of that, than anyone else. And if you want to take risks, I think you should risk... I'm actually fairly uncomfortable with some of the higher risk philanthropy, sort of by default or generally and I feel that when you start doing this kind of giving with a 10 or 20 year time frame, you're doing something inherently very difficult and it's easy to be overconfident, it's easy to fail. And so, I want to be up for that by, if you can win, I want the win to be really, really, really huge.

0:57:06 HK: And so the kind of stuff we do at Open Phil, I like it and I'm excited about it because I think the wins are big enough to justify that. I frankly think a lot of philanthropy falls right in the middle where people are doing things that they're on a 10 or 20 year quest to change some giant social problem, but their best case is actually only okay and even their best case makes it kind of hard to come up looking better than Give Well's top charities. So for me, I like both extremes, quite a bit. I personally give to Give Well's top charities, so I believe in allocation to both. I think some of the... If I look at the... Maybe not surprisingly, Open Phil pursues a little bit more high variance, so I feel that the variance Open Phil is doing, I like better than Give Well's top charities, but I wouldn't say that about all of it and I continue to have unstable views about this. I personally have recommended both and give to both and certainly haven't just dropped one from my portfolio. I can take one more.

0:58:08 S?: Can you give an overview [0:58:10] ____ by cause area? This is for animal welfare and SAI [0:58:15] ____ budget breakdown [0:58:17] ____ and [0:58:18] ____ budget by cause area.

0:58:24 HK: Sure. So how does the budget breakdown by cause area and how do you decide that? So far, we've not been very consoled about it. So far, everything we've said is we want to start giving. We want to start experimenting and when we're in a cause, so there's two different paradigms of giving, working on cause. One of them would be, we have a very casual interest in the cause and someone who's a generalist on staff would work on it. And in that case, it would just have to be by the grant and someone finds a good grant, we make it, we think that grant is better than the last dollar, or it's a good learning opportunity. And then when we hire specialists, it's more important to set a budget because we might be hiring somebody really passionate about farm animal welfare. So if we make it all about the grant that's going to give them too much incentive to play off the goodness of their grants and everything else, then we need to set a target.

0:59:15 HK: Again I want those targets to be more principled in the future than they are now. For now, a lot of what we have done is we just said, First off, we want to make it big enough commitment to this cause that we can be a major player, get top talent, be reliable for the field and from major commitment. We've kind of looked to what other foundations major program areas looked like and it tends to be in tens of billions of dollars per year. And then we've made adjustments from the starting point, where we've had to have an annual portfolio review process. Where the program officer would say, "Well, if I had a little bit more money, this is the kind of thing, I would do more of and if I had a lot more money, this is the kind of thing I would do more of." We look at it and look for some sort of inflection point of diminishing returns and go with what feels good but it has not been very principled, so far, so for [1:00:01] ____ choosing right causes. Choosing the right people, having the right process. But getting the right dollar amounts, that's a future thing for us, and it's in progress. Do you have a follow up?

1:00:10 S?: Yes, are all of them funded roughly the same amount? Or.

1:00:14 HK: No, it did vary by cause. With the special causes it's going to be tens of millions per year. And with other causes, it's going to be by the grant so they will vary a lot and it just depends on just what the year was. Cool. So that's all the questions I'll take for now. I'll stick around for 10 or 20 minutes if you folks just want to come up and ask more questions. And thanks everyone for coming.