

# GiveWell San Francisco Research Event June 6, 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.*

0:00:02 Holden: Alright, everyone. I'm Holden. I'm the Executive Director of the Open Philanthropy Project. I've just turned the recorders back on, so I think you already got this message, but if you don't want what you say to be public, just say so when you say it or tell us afterward. And otherwise, this is all going to be... I think the transcript goes up online. So I am going to talk about the Open Philanthropy Project which, as of very recently, is actually a separate organization from GiveWell, so we completed that process. We're going to have a public post in the future describing why we did this, what some of the details were, but nothing really changes overnight. We're still in the same office and de facto I had been at GiveWell but spending basically all my time on the Open Philanthropy Project. So it's the same situation conceptually but it is a different situation formally, and we do ask that people not tweet that or whatever, because we haven't quite gotten our public announcement up, but it should be up in a couple weeks.

0:01:03 Holden: So the Open Philanthropy Project began as GiveWell Labs in 2012, but I think it's worth talking about the differences and why we've ended up becoming separate organizations. Basically the Open Philanthropy Project is, like GiveWell, trying to find the best ways to accomplish good by giving and trying to share a lot about what we learn to help other people decide what to do, but we do things differently, so our main sort of target audience is major donors. Especially, right now we're mostly trying to help Cari Tuna and Dustin Moskovitz give away their fortune, and I think at the point where we're closer to that goal, we may have more of an emphasis on also trying to help other donors but they certainly have a lot of giving to do. We spent the first few years since 2012 really building a framework, choosing the areas we were going to focus in, and it's only last year, 2016 was the first time that we really got giving up to a high level. There was over \$100 million in 2016 that was given basically at the Open Philanthropy Project recommendation. That does not include any of the Cari and Dustin gifts to GiveWell top charities. It doesn't include anything to GiveWell top charities, it was all separate things, which I'm going to talk about.

0:02:24 Holden: And so, where we are in our development is we have just kind of reached a point where we're not at the level of giving we want to be eventually, but we're at a pretty high level, and I think for the next couple years we're going to be trying to just build the best possible organization and process around that giving before we pursue another round of big scale-up. So in terms of what we do and how we do it, we spent, like I said, our first few years building a framework and choosing causes, and we looked for causes that combined being important, neglected and tractable. And instead of walking through exactly how that process went and why we chose what we did, which you can read about on our website, I'm instead just going to emphasize a few themes of our work that have kind of emerged as we've tried to find the most important, neglected and tractable things. So first off, one consistent theme of all the Open Philanthropy Project work is that a lot of it is what we call hits-based giving.

0:03:23 Holden: This is something we wrote a blog post called Hits-based Giving, explaining the concept, but it's basically kind of a similar idea to venture capital in a way that you might make 10 bets or 10 grants or enter 10 different causes and be totally fine with failing to have any impact in nine of them, and with the tenth one, which you don't know in advance which one it will be, having enough impact that it kind of makes up for everything else and then some.

0:03:52 Holden: And so this is probably a familiar model to anyone who's in Silicon Valley and reading about Y Combinator. I think it also is somewhat familiar from the history of philanthropy so there have been some huge successes that I think are attributable to philanthropy. A couple of examples that are particularly huge and we have conference rooms named after some of them, but one of them is the Green Revolution, where the Rockefeller Foundation funded this improved agricultural approach that a lot of people have credited with sort of kick-starting the East Asian Tigers development phenomenon, getting a lot of countries to go from poor to middle-income, turning India from being in the middle of a famine to being a wheat exporter, and saving over a billion people from starvation, and resulting in a Nobel Peace Prize for Norman Borlaug. So that's like a big win. Maybe...

[laughter]

0:04:40 Holden: Maybe one of the most significant humanitarian developments of the last century, or really ever, and it really did start with the Rockefeller Foundation funding some research that I think if you had asked them to prove to you that it was going to work, they... If they had had a proof, the proof would have been wrong. And another good example is the pill, the most common contraceptive, was a case of Katharine McCormick, a philanthropist funding something that society I think was not necessarily really ready for or on board with. And certainly, that government

wasn't going to fund. And it was even a case where they, I think, were not able to advertise it for its main purpose of preventing birth. They had to advertise it as kind of a weird fertility treatment but they got to put a warning label on there that you couldn't actually get pregnant while you were taking the pill. They were required to put that on there and that kind of served as their advertising.

0:05:36 Holden: But those were examples of where I think you'll see a different philosophy from the proven cost-effectiveness framework where there are cases where a philanthropist has been able to do something that was really kind of out there and different and really had a good chance of just totally failing and probably for those two successes and many more like them, there's lots of failures for each success, I would be sure. But it may all be worth it in the end, and that's a lot of the approach that we try to embrace, and it leads us to do things pretty differently. So, we have tried to pick causes that are a good fit for this sort of thing, and with that in mind we've really tried to pick areas to work in where we can just sort of in a sense have the highest potential to hit homeruns that one could imagine.

0:06:26 Holden: If there was a way to do just an enormous amount of good that people would look back and say, "It would have been worth failing 10 times if I could get that." What would they be? And so two major themes that I would say cover most of the causes we work on, two major kind of theories of how a cause could be good. One of them is the idea that I call radical empathy. We have a blog post laying out the idea, but basically over the course of history people have kind of gone from only caring about their family, or their nation state, or their tribe to only caring about their country, to at some point recognizing civil rights and gender equality. And these things seem like among the bigger pieces of human progress ever. And if you look back and you say, "How could I have done a really huge amount of good?" I think a good way to do that would've been being ahead of the curve, being early on that stuff, and sort of recognizing people as your fellow persons when others didn't.

0:07:28 Holden: And so a lot of the work we do is trying to serve or help people who are not necessarily considered to be in the inner most circle of concern by most of the world. Certainly global poverty fits in there, certainly the GiveWell top charities which now Open Philanthropy is in the position of deciding how much to recommend to those top charities from Cari and Dustin. Certainly that's in that category because I think that when you look at who Americans are prioritizing and how, definitely how the national conversation sometimes goes around immigration policy and things like that. A lot of times people don't feel that overseas people, people in Africa should be prioritized, and we feel they should be. Other examples, we do work on criminal justice reform which we think it's an unusually tractable political area, one where we see more opportunity to sort of get some wins than elsewhere. And that's another area where a lot of the justification, a lot of the reason for that cause is that we think there's a huge amount of harm done to people who are in prison. Often not at the very top of the priority list for people having debates about criminal justice policy.

0:08:37 Holden: And then something that's maybe a further step out there is farm animal welfare. So this is a case where there's really, really large numbers of animals being treated really, really horribly on factory farms. And we've, I believe, learned that there are some really compelling tractable opportunities to help. A lot of times, it's as simple as running a campaign against a fast-food company or a grocer asking them to avoid embarrassment by making a pledge to treat the animals a little bit better. And that can make a huge difference, and the thing is, a lot of people would look at that situation and say, "Great, you made a huge difference for chickens, who cares?" And we would say, "If 50 years from now we look back and we kind of wish we had cared about chickens sooner, this would look like one of those really big hits."

0:09:22 Holden: Other major theme that a lot of our causes fall into, is what I sometimes call X factors for the future of humanity or maybe more specifically, global catastrophic risks. So that's the idea that a lot of the best things that have happened to date, a lot of the biggest reductions in poverty haven't looked so much like money getting redistributed from rich to poor, as they have from just massive transformational changes that swept along all of the world at one time. And so the Industrial Revolution being really the big one, that I think you would credit with much more of a drop in poverty than any kind of specifically anti-poverty intervention. And so the question is: Could there be transitions? Could there be sort of things that change the whole nature of civilization coming up in our future, that we might have an impact on for better or for worse? And they could be good things that we want to speed up, they could be bad things that we want to prevent, they could be things that could go either way.

0:10:20 Holden: But certainly if you were looking again for a way to look back in the past and wish you'd done something, you might wish you had played some role in the specifics of how the industrial revolution played out. And looking forward, with the world as kind of interconnected as it is, you could definitely imagine things changing something about the whole planet at once, and so an obvious example of this would be climate change. We've funded work on reducing greenhouse gases, also on preparing for worst-case scenarios where we might need to deploy solar

radiation management to counteract the effects of climate change. That could be a good or a bad thing.

0:10:56 Holden: And then there's a couple other global catastrophic risks that we've prioritized even more highly because we think they not only have huge potential to be very big risks but that they're very neglected. So one of them is biosecurity and pandemic preparedness. So if you kind of made me look into a crystal ball and I saw that humanity got totally wiped out or a large chunk of us got wiped out in the next 100 years and I had to guess how, I would probably guess that it was some sort of pandemic, perhaps a synthetic one. Biology is advancing in a way that I think it's going to get unfortunately easier to design new kinds of pandemics. And so I think the world is underprepared for that, I think especially the biosecurity community does not take very seriously the idea of really massive pandemics that could spread across the whole globe. I think they tend to focus a little bit more on the probability of tragic but more contained pandemics; and so that's an area that we work on.

0:11:53 Holden: And then potential risks through advanced artificial intelligence, which would definitely be a whole can of worms to go into too much, and I want to get to questions. I may go into more detail later but again, we think one of the most dynamic and unpredictable areas of science right now that could change the world a lot, is artificial intelligence work, and we can envision ways in which that could go really well or really poorly. We think that some of the risk management research that can be done is the most neglected in the field, and so we are trying to build a field around people who are doing research to reduce the risks of AI. So those are the causes we work on, those are the areas, that's not all of them, those are the main ones. One more thing I'm going to talk about before I take questions, is just the basic MO because a huge part of who we are is defined by what causes we work on. And we think we've put more thought and attention into picking the best causes to work on than other foundations have, and that's a huge distinguisher for us.

0:12:51 Holden: But once we pick 'em we also need to figure out where to actually give the money, and there I think we also have a somewhat, although maybe less unusual model. The basic MO, and again this will be very different from GiveWell, we work really hard to hire the best people we can for the role of program officer, and the role of program officer basically becomes the person who takes the lead on all of our strategizing and all of our grant making in an area. So they become the point person and they are... Everything we do is because they suggested it. Now, we question them a lot and we especially me and Cari.

0:13:26 Holden: So Cari Tuna and I; Cari's the president, we are the ones who give final approval to grants. And a big part of our process is when a program officer brings us a proposal, we will ask a lot of critical questions, we have internal write-ups that we've shared versions of online, that ask for things like what predictions can be made, what do we think the outcome of this grant will be, with what probability, how good would that be for the cause we're working on, what are some concerns, what are some ways it could go wrong? And so we definitely spend a lot of time syncing up with our program officers and trying to make sure we understand what they're bringing us, but in terms of leading the strategy, in terms of deciding what is worth thinking about, what is worth contemplating and what isn't, that is really in the program officer's hands. And we've designed it that very way, very deliberately because we think that gives us maximum flexibility.

0:14:14 Holden: So we aren't stuck with a process where your grant needs to check these five boxes or we can't look at it no matter how good it is. Our process really, we have to get the right person and we have to sync up with them, and then we're able to do anything that, in their judgement, is good. And that plays to, I think, our strength as a philanthropist that we're not a big bureaucratic agency, we're not a government, we can do things that are different, we can do things that are risky, we can do things that are out there. So a final thing that kind of characterizes this process is we have this rule that we call the 50/40/10 rule, which is that each program officer should be aiming for about 50% of their portfolio by dollars to be grants that the decision makers, Cari and I, really kind of agree with, that we understand the case and having heard the case, we would make the grant if it was our money. It's Cari's money.

[laughter]

0:15:08 Holden: We feel affirmatively convinced and that we could personally defend that case. Another 40% of that portfolio, so bringing the total to 90%, should be things that I look at it and I say, "I could easily imagine how I could see the case for this. I don't feel totally convinced but I can totally see how if I knew more and if I was in your shoes and if I was the program officer, I would totally get it." And then the other 10% is basically discretionary, so we would say, "This grant does not make sense to me. I don't understand why you think it's good, we're making it anyway." And that setup is supposed to make things such that the program officer does need to convince us of things and we do need to be 50% synced up and 50% in agreement. But it also means that they can do the things that make sense to them and

we don't have the whole organization relying on things that can be explained to non-experts, which Cari and I are.

0:16:03 Holden: So we're trying to really put the decision-making power and a lot of the strategizing power in the hands of the people who know the most, and I think that positions us to be flexible and to do things that are just a little bit different, a little bit risky and have the best chance of being that kind of hit. So that is the causes and the overall MO and I'll take questions for a while now. Yeah.

0:16:26 S?: I'm curious to hear more about what you guys are doing around criminal justice reform. You probably came across the Slate Star Codex write-up recently on bail assistance, for example. I guess there is a couple of different ideas about what to do in that area.

0:16:38 Holden: Sure, so on criminal justice reform... So the question is: What's our strategy on criminal justice reform? What do we think of bail payments for individuals? We have supported some work on bail funds. It's been more, again with the hits-based approach, it's been more trying to create bail funds that can then sort of sustain themselves or can bring in money from others. And that's been more the value proposition and also trying to, in many cases, make the case that money bail is just not a good system, and we'd do better off without it, and it could be a way of creating a lot of benefits to get rid of it or to dial it back a lot.

0:17:17 Holden: So what are we doing in criminal justice reform? Certainly Chloe Cockburn is our program officer there; she came to us from the ACLU. She is the one who sets the strategy and so that is the highest order answer. And then in terms of what we're specifically doing, one of our big grantees is the Alliance for Safety and Justice, so they grew out of Californians for Safety and Justice. Californians for Safety and Justice was behind Prop 47 and Prop 57, two major recent initiatives that we feel are going to be very impactful in California and were very well run initiatives.

0:17:50 Holden: And so Alliance for Safety and justice is like the new national version of that. We were, I think one of the biggest or maybe the biggest initial founding donors to that effort. And they are basically going for sentencing reform so they're trying to basically make the case and pull together advocates and coalitions to get bills passed, and they had a big win, we believe that we're checking out at the moment to see how big a win it was, in Illinois, with a bipartisan piece of legislation to reduce incarceration. So that's been one of our major bets. We've also been funding a campaign to close Rikers Jail in New York City, and there was a verbal statement of intent to close Rikers Jail following the efforts of that campaign, which we think is great. And just to be clear, closing the jail would require cutting the jail population. It would not just be a matter of moving people around. That's why we're interested in it. And some other...

0:18:49 Holden: So that's a lot of sentencing reform and just going directly for the win kind of thing. We've also supported some work on exploring alternatives to incarceration. So just things that fulfill a lot of the same functions of holding people accountable, that don't necessarily cost as much in financial or human cost. And so, supporting some pilots and advocacy on those and then a major effort has also been prosecutors. So, we believe that in addition to the laws, a lot of the reason incarceration in the US is as high as it is, is because there's a culture among prosecutors, where there's kind of a belief that it's the more people you lock up the longer, that makes you tougher, that makes you better. And kind of challenging that and making the case, that the real job of a prosecutor is to think about all the interests of their community and to not overincarcerate people, because the people in their community are affected by that financially and in a human way. That's been a major priority for us. Those have been the big pillars of what we've done so far.

0:19:50 S?: Could you see a time where you would go narrower and deeper? I'm impressed with the breadth. But it's also like, "Oh my gosh. That's a lot."

0:19:58 Holden: Sure. Yeah, could be go narrower and deeper? Are we spread too thin?

0:20:01 S?: No, I wasn't saying you are.

0:20:02 Holden: Sure, yeah.

0:20:02 S?: Does it look like, look into your crystal ball. Could you see that happening? Look at that...

0:20:08 Holden: Well, the thing that I could definitely see happening is that we don't add much more and that giving... We still need to bring annual giving up a fair amount. And so I could imagine that what it might look like is that we stay

in roughly the causes we're in but that giving goes up a lot and that commensurate with that, we hire more people in each cause. And so there's more people working on each thing. There's more depth. There's more expertise. There's more dollars. That could be a direction we go, rather than just adding more and more causes like this. I think that's a strong possibility. I could also imagine deciding some day that we think we've gotten better results in causes we've understood better and that we believe we have an opinion on which ones are the best ones to focus on. And so we do want to cut some and focus more, but I think that's somewhat unlikely. I do think that a lot of how we're positioned as a funder is that our number one priority is definitely to support other people to do great work and not to make everything happen ourselves. And so sometimes there's definitely elements for us of making the case of advocating for ideas, not just of writing checks.

0:21:17 Holden: And I can get into that more if someone wants. But, in general, the end goal of any project we embark on, the end goal of anything we do is for someone outside the organization to be the real expert and for them to get funded by us and for them to do the work. And a lot of times we approach grantees in a similar spirit to how I've described our relationship with program officers, which is they need to convince us enough that we understand, that we believe in them but we also leave a lot of the discretion to them. I think in a lot of ways it is appropriate for us to be pretty broad. A lot of foundations of comparable size have a lot of causes. And I think that rather than seeking to become the absolute experts in a cause, I think we want to be the meta experts in identifying who the experts are and supporting them. And to do that, it doesn't necessarily take a huge staff in each area.

0:22:07 S?: Can you talk a little more about your radical empathy portfolio?

0:22:10 Holden: Sure. The radical empathy portfolio. Radical empathy is not a cause, it's more of a theme that describes a lot of our causes. The things that fall under that, I listed global poverty. I listed criminal justice reform, and I've talked about criminal justice reform and, I think Cat talked about global poverty in a sense because the main thing we do on that is supporting GiveWell's top charities. And that leaves farm animal welfare, and I gave the high level picture there, but Lewis Bollard is our program officer there, and I think that's the area where we've had the most short-term, immediate tangible success. It just feels like there were a lot of victories that were ripe to be had by pressuring corporations to make pledges. And I think in a lot of ways, corporations are easier targets than governments for advocacy. And so I think it's been pretty interesting to watch that. And Lewis put up a blog post pretty recently about why we think it's been so successful.

0:23:08 Holden: To be clear, the successes belong to the groups we're funding, and I think a lot of that was already on the way when we came in and we've tried to support it, accelerate it, speed it up, broaden it, but I do think there have been successes and I think it's been pretty interesting to watch. One of things we did in that field, is we came in, and we immediately scaled up, funded US corporate campaigning and then that, very quickly afterward just got very scaled up wins and kind of got all the major grocers and all the major fast food companies in the US to pledge cage-free. And it happened in the year after we poured in that funding. And I think some of that was inevitable and some of that was already on the way for previous victories. And then what we've done since then is really try to think, "How far can we push this model into other parts of the world?" And so, we've spent a lot of time trying to figure out, not only Latin America and Europe, about how can we support farm animal welfare work in China and India, which are going to be huge sources of meat consumption and growing as they get richer. In the back corner.

0:24:16 S?: How are you going to compare the impact and success across the projects?

0:24:22 Holden: Sure. How do we compare impact and success across projects? Evaluating our past impact is something that we have not done as much of as we eventually need to. Partly because last year was the first year when we really did a lot of giving, and so I think the first year that it makes sense to look back and really evaluate ourselves is probably going to be about two years from then. So, it might be a year from now. And the answer is kind of a work-in-progress. I can give you some previews. Whenever somebody puts through a grant, the grant process is that someone completes an internal write-up and they answer all these questions, it's in a Google Doc, they share it, and then Cari and I will ask questions and give it a thumbs up or not, but I usually end up approving. And those questions include things like, "What are you hoping for from this grant?", "What would success look like?", "How would we decide whether to renew?" And then it also includes a section for probabilistic forecast, so it says, "Can you make some tangible measurable predictions that are related to this grant even if only vaguely related and can you put probabilities on them?"

0:25:25 Holden: For example, if we're funding a group to scale up its campaigning in India, we might say, "I think there's a 80% chance that at least one corporation is going to make a major pledge on animal welfare in the next two years. I think there's a 20% chance that a large number of corporations will make a pledge," that's quite impressive. And

sometimes the words will be that, there will be some subjectivity and so the idea is for us is to go back in a few years and score them. And I think scoring them will take some judgement, but at least then we'll be able to say, "Okay, when we said something was 80% likely to happen, did it happen 80% of the time?" And that's one way of checking ourselves.

0:26:04 S?: Do you ever consider political action and do you think it will create progress in one of these areas, or do you try, as a separate goal, to progress but remaining apolitical?

0:26:14 Holden: We don't have a particular, like in a broad sense, we don't have a particular goal of remaining nonpartisan or anything like that. "Political" has a certain meaning that I don't want things to get confused by, so I'm going to talk about maybe instead of like partisanship. We're not trying to sit in the middle and be non-judgmental and offer technocratic advice. A lot of the work we do just takes a very definite side. It's a side that we think is best for the world and then we will push it. The organizing idea of the Open Philanthropy Project is do as much good as possible, and with the hits-based giving idea we don't want to shy away from things just because they might be controversial, just because we might make enemies, just because people disagree with us. And certainly in criminal justice reform, we have a point of view on that, and it's a point of view that we've tried to do our homework on, a lot.

0:27:10 Holden: David Roodman, who is our best critical evidence reviewer and the best one I know of, generally, spent a year going through all the best papers that you could use to get a sense of the effect of incarceration on crime, and ended up concluding that at the present margin, the best guess for where we are in America, that incarceration is giving us a zero benefit in terms of public safety. That it is, best guess, as criminogenic as it is crime-preventative because there are ways in which there's more crime that can result from people going to prison. And so certainly a lot of research was done there. He came in as someone who didn't really have a dog in that fight, has very meticulously gone over the studies, has tried to do it in a very objective and not prejudging way. But early in the report that was where it looked like the conclusion was likely to end up, and that was our best guess of how it would end up based on our understanding of what the top experts in the area think and some other things. And so having that side, that is what we push for. We don't just try to advise people. We do say, for example, money bail is a bad system and we have an agenda of trying to make it more likely that it goes away.

0:28:34 S?: When we talk about global catastrophic risks, for example, the refugee crisis, the global refugee crises, do you have any sense that this is something that you can find a handle on or have a role in? Or...

0:28:49 Holden: Sure. So the global refugee crisis, we have done some work on immigration policy. The priority of it has fallen a little bit for us, but again, that's a place where we have a perspective, we have a side, pro-immigration basically, especially from the poorest countries. But as far as global catastrophic risks, that term has a specific meaning for us that, it's about, what we consider global catastrophic risk is something that could change the whole global nature of civilization in a way comparable for better or for worse to the industrial revolution, and so that is a place... I would call the refugee crisis, I would call it humanitarian crisis. I would call it a tragedy, but not a global catastrophic risk by our definition. We're looking for things where there might not be humans anymore or the whole world order might work in a different way, and those are the things that fall in that category. This one might fall more under the radical empathy theme.

0:29:45 S?: What sort of prediction goals do you find yourself setting on global catastrophic risk projects?

0:29:49 Holden: Sure. So predictions are especially hard to make on global catastrophic risks.

[laughter]

0:29:54 Holden: I will maybe go into a little bit of a digression on AI, I imagine that's something some people would be curious about, and it can give a good sense of how we go about this. So I think artificial intelligence research is an exciting and dynamic field right now. I think most people who work on it, they go into work every day saying, "How can we get AIs to be able to do more things? How can we be able to get them to reason better, to be smarter, to take on more of the tasks that humans can do?" And there's another set of research questions that says, "Well, what are some of the risks here?" And one example risk is that common AI paradigms right now, the way people do AI research, is often very reliant on having a very well-defined reward function, and so if you have a system that will sort of make a guess, and then if it's wrong, it gets sort of an algorithmic punishment. If it's right, it gets a sort of reward. And from that process, doing it over and over again, it learns how to get what it's going for. And the problem is that if you imagine AIs getting much more broad and much more useful, there are certain goals that are very easy to define, like, "Maximize the

amount of money in this bank account." And then there are other goals that are very hard to define and that we have only the fuzziest sense of what they even mean, which is like, "Make the world a nice place."

0:31:16 Holden: And the problem is that generally the things we care most about and the things we want to do, are the things that are harder to define. And so there's a concern that you could end up with a future world in which you have these very powerful algorithms that are very smart and very effective if you ask them to maximize the amount of money in a bank account, but no one really knows how to get them to accomplish something fuzzier, like for example, stop other AIs from messing things up when they're maximizing the money in bank accounts. And "messing things up" being a vague term that I think we don't fully understand ourselves what we mean by that.

0:31:47 Holden: And so there is research one can do on trying to get ahead of that kind of problem. So an example of that kind of research would be AIs that are able to learn fuzzy, poorly defined goals from human feedback. So an example is a recent paper with our technical advisors as main authors that basically, there's a simulated bot that usually people try to teach it how to walk by something like teaching it to move its limbs in the way that it moves forward at the fastest possible rate. And in this paper, they try to teach it to do a back flip, which there wasn't a ready to go reward function for. And they try to teach it that by kind of just, it would move randomly and a human would watch two different films of it moving and say, "That one looks more like..." kind of playing hot and cold. "That one looks more like in the right direction." And so from that it would learn to do a back flip.

0:32:39 Holden: And so the high level point here is that there's a certain kind of AI research that we wish people were doing more of to get ahead of certain risks. And there it's kind of a complicated thing to try and get that to happen. And so the ultimate goal is reduce some very small, very far off risks. That there's some huge catastrophe that happens from very intelligent agents trying to maximize a very narrow and unhelpful goal. But there's an intermediate goal, which is, we would like there to be more really good people working on, what we consider sort of safety-biased research, or catastrophe prevention research. And that is sort of measurable, in the sense that we can say right now, "How many people are working on this?" and we can come back in two years and we can say, "How many people are working on this?" And then we can ask our technical advisors, the people we trust most, "How good are these people? How good is this research?" And we can write that down and try to make it go up. And there's a lot of subjectivity there, but that is the kind of thing we'd be making predictions about is: How many people will be working on this problem? How many papers will come out that we think are good? And that's an intermediate metric. Yep.

0:33:49 S?: [0:33:49] \_\_\_\_ history of field growth posted awhile ago. It seemed kind of light on, "Here are the specific recommendations we have from this field growth." Do you have [0:34:00] \_\_\_\_ on that or do you want to say something about how it relates to growing that field of AI safety?

0:34:06 Holden: Sure. So we recently put out a report on philanthropic attempts to build fields. And so the report kind of talked about a lot of fields. One in some depth, and the others in not so much depth, that philanthropists have intentionally tried to build, and this is relevant to our work on AI. It's also relevant to our work on other things, but it's relevant to this whole idea of... Currently no one works on this. We would like lots of people to work on it. What should we do? That's not an obvious thing to solve. And yeah, the report was pretty light on details on what to do because it's very hard. In general, when you run across success stories of a philanthropist helping to build a field, what you find is that they were sort of funding everything and just doing everything that might work. And then the field grew. It's really hard to look at that and say, "Why did this happen?"

0:34:56 Holden: And I'll give another example from AI, which is that there's often a chicken and egg problem, and certainly we perceive there to be one in AI, where it's hard. Senior researchers, they have their interests. They know what they're interested in, and so it's not going to be that easy to convince them, or make the case for them, or interest them in switching fields to some other topic they haven't spent most of their life thinking about. On the other hand, if you try and get junior researchers to work on a certain topic, it's going to be hard to get them to work on things if they don't see senior researchers working on them, 'cause they might then be skeptical that there isn't much of a career path there. And so, what we've ended up doing is we've said, well, we're going to try and fund whatever senior researchers we can, and then we're also going to try and have a fellowship program, which we're working on, for junior researchers. And the two are just going to have to play off each other. I kind of don't fully expect one to work that great without the other. And so it's going to be very hard to point to one and say, "This is the thing that worked."

0:35:53 Holden: And so that I think is why that report was light on detail, is just that the main thing we learned is that if you're looking for times when philanthropists went and tried to build a field, and spent a ton of money, and then the field got really big. Yeah, there's lots of those cases. If you try to figure out exactly what happened and why, you can't

really figure it out. So our take away was, we may as well give this a shot 'cause there's plenty of precedent for it working. Luke, do you want to add anything to that?

0:36:17 S?: No, it sounds good.

0:36:18 Holden: Cool. Luke wrote the report so you're going to give it a shot? [chuckle] Yeah, in the back.

0:36:25 S?: So [0:36:26] \_\_\_\_ ask about uncertainty and risk, and then time value of money, or time. So one view is that GiveWell says, "You give to this charity, you give 100% to this one charity [0:36:40] \_\_\_\_ you're pretty certain about that," and they're saying that. And if I'm a small donor, do I just follow that or if I have more money, like \$100 million, a billion dollars, "Well, I should open up a Open Philanthropy Project [0:36:51] \_\_\_\_ understand these areas, I'm going to wait for those outcome, they're uncertain." At what scale would you say you should switch from just give to this and do that, or we're early in this stage or Rockefeller, everybody has been doing this for hundreds of years, we're just one other guy, we're not going to... We want a better recommendation in five years, so just do whatever. Or, "Hey gosh, this is really new and exciting, in five years there'll probably be 10X return on investment if you wait."

0:37:21 Holden: Yeah. Sure, so the question is how should a donor decide if they want to, for example, do this hits-based giving that's very long-term and very ambitious versus if they want to do more of the linear metrics-based, "I can see my impact and see how I'm doing each year that I do it." I don't have a full answer to this. I think one answer is that we wanted to try the thing that wasn't already being tried in the way we wanted to try it. So first, there was GiveWell and then there was Open Phil. In a sense we were always looking to do the thing we had an opportunity to do, that we didn't already see being done.

0:38:01 Holden: I think GiveWell has been thinking about putting out, or maybe they did put out, but they've at least been thinking about putting out a post on this kind of question, on like, "How do you decide which kind of donor you are and what kind of thing you should do?" But I think one really important access is not so much about how much money you have, it's about maybe how much time you have. And also how much trust you have. So I think, in my opinion, the best way to do the hits-based giving thing is to either get really deep into an area and be an expert in it, or find someone who's really deep in the area who's an expert who you can either trust, or build trust with. Because in my opinion, to do really good hits-based giving, you want to do things that no-one else will.

0:38:41 Holden: And so, if the cases can be laid out in public and it's all clear, and all the evidence is there, and it's worked 10 times before and you want to do it an 11th time, it's pretty unlikely that that's going to be the kind of thing that you're way ahead of the curve, and no one else has seen it. And so I think my general framework is that people who are really deep in the details, or who trust people who are, are the best position to kind of see something that other people are not ready to see yet or aren't able to see that is very risky but that there's a good argument for trying.

0:39:16 Holden: And I think if you try and engage in this hits-based giving without that kind of depth, I think you'll probably just run into things that look exciting which is almost the opposite of what one wants, because those are the things that people will snap up. So that's my basic guideline and I think for someone who... They work all year, and then they have a few hours each year to decide where to give to charity, I think GiveWell makes more sense because I think for GiveWell you can say, "Okay, I understand the case. The case is being made, my money will do a lot of good here." And I think you do need either the expertise or the trust in the expertise to, in my opinion, do the other one really well.

0:39:52 S?: Is the expected value or the history of hits-based giving greater than...

0:40:00 Holden: Yeah.

0:40:00 S?: The more metric-based simple one or...

0:40:02 Holden: Right.

0:40:04 S?: In theory I would expect it must be otherwise you wouldn't have opened it, right?

0:40:07 Holden: Oh no. [chuckle] So one question is which one is better and which one looks historically better, hits-based giving or the more linear approach. And I don't know which one is better, and that's why I started my answer the way it did. 'Cause a lot of what we wanted to do was just try it. We wanted to try hits-based giving in a way that we

could track how it goes and in a way that we could form a better opinion of it from being in it. So I think this is unknown. I think there are enough examples of big hits that I think hits-based giving is worth a shot. But I don't think it's tractable to actually figure out the return on investment. It's too hard to find all the things that didn't work. Yep?

0:40:40 S?: I got a little confused there. Is hits-based giving really not a linear phenomenon? I imagined it was essentially expected value-based giving, where you're willing to pick on a very large amount of variants, and then that combined with you hiring your own experts who you believe are able to make better unknown estimates perspective value, rather than relying on public information. But I didn't think...

0:41:00 Holden: Sure.

0:41:00 S?: Is that what it is?

0:41:01 Holden: Sure, yeah. That is what it is. This is just a clarification on the definition. When I said "linear" I was just trying to, more, evoke an image of kind of like you move a step at a time. And so you have this model in your head that if you give \$3,000, something good will happen, if you give another \$3,000, another good thing will happen. This sort of model has proven out, and I think with hits-based giving it's a lot more of like, "Here's an opportunity. It looks totally different from all the other opportunities. It might work." Now we tried funding that, now we're looking at something totally different. But I don't think we need to get hung up on the term "linear." I think you're correctly describing it. Yeah.

0:41:34 S?: I also think if you want to look at other ways of thinking about it, even GiveWell, in the presentation we saw before, they said by their best estimate it was only a one in five chance that deworming the world were even working, so that's already... It may or may not work.

0:41:47 Holden: Sure.

0:41:48 S?: But even if you think Deworm The World works probably you give money and some large fraction of the kids you deworm, didn't have worms anyway so they didn't need it. So it's not linear in that way, so I think the risk occurs at all levels. I think it's just a question of how broad you want to spread your portfolio and what you want your risk to be.

0:42:03 Holden: Sure. I think it's totally fair to just describe it as a spectrum and to say that the hits-based giving is sort of more one extreme, and the GiveWell approach is more at the other extreme. And certainly with GiveWell, there's a lot of possibilities that things will not go the way you hope. But yeah, I think you have a decent picture, and I think people feel fairly clear on what the contrast is so I don't think we need to get all the terms exactly right. Yep?

0:42:29 S?: So hits-based giving it sounds like it's heavily reliant on expertise of the program officers. So how do you find and vet those people and trust that expertise?

0:42:39 Holden: Sure. So hits-based giving is dependent on the expertise of the program officers. How do we find and vet them? So the basic answer there is that we try to learn enough about a cause that we can do some meaningful vetting of people. And so we try and get into it ourselves enough to be able to do that. And so with criminal justice reform we talked to a lot of people who've made some exploratory grants, we heard a lot of perspectives, and by the time we were interviewing people we kind of had this whole list of things in our head that people had said would be a good idea to try.

0:43:09 Holden: And so if we're interviewing someone... A lot of the people we interviewed, they would know a huge amount about one part of criminal justice reform, about one sort of corner of it, but then when we questioned them about other stuff we'd heard, we often felt that it wasn't really something that they... It didn't really feel like they knew more than us about all the different areas and they could educate us on all of them and they could kind of choose between all of them in a way that we would if we knew more.

0:43:36 Holden: And so that was a major part of how we did that is we just tried to have a lot of stuff in our heads, and then we tried to find the person who felt the most like they were just a better version of us, and they kind of answered everything as if they had been us but thought a lot more about it. That was kind of the idea we were going for. And knew a lot more. And that's also how we assess and so it's definitely a spot-checking process; it's definitely like, the grant writer comes in, and I'll take the parts that make the least sense to me or seem the most interesting to me and I'll

ask a lot of questions.

0:44:09 Holden: And what I'm really looking for is that it feels like this is just like, this person is more expert and is better, and it's not that we have different values or we disagree or that there's areas they've never thought about, it just feels like they're doing a better job than I would. And so that's the basic MO. We have a blog post on this that I think is on our Notable Lessons page which is like, "How we hired our first program officer," and that is the method we've used for basically all of them. Yeah?

0:44:37 S?: It seems like there's at least two levels of risk management that you have to do. There's one among all the topics, you pick a particular topic like animal welfare, you have some estimate of how risky it is that this is the right thing, that this will be valued in the future was the way you put it, right?

0:44:55 Holden: Yeah.

0:44:56 S?: Versus the probability of succeeding and all that, right?

0:44:58 Holden: Yep.

0:44:58 S?: And then within that topic, there's a spectrum of very likely to work, keeping it real kind of approaches like pressuring fast-food restaurants, right?

0:45:11 Holden: Yep.

0:45:11 S?: Where you can envision all the way down the spectrum of doing research into cell-cultured meat growth things, right?

0:45:18 Holden: Yep.

0:45:19 S?: Or if you think that's likely to happen, then maybe the value of pressuring fast food restaurants goes down anyway, right?

0:45:25 Holden: Sure.

0:45:26 S?: Do you modulate the risk within a category based on the likeliness of the overall category or how does that work out?

0:45:36 Holden: Sure. So how do we choose... So there's the risk of like the overall theory of the cause, maybe we'll be glad that we decided to prioritize the welfare of farm animals, and then there's like a whole spectrum of risk within a cause like are you aiming for a 10-year strategy to get one giant win? Or are you doing something that's more concrete like maybe bail reform or some of these corporate campaigns? And the basic philosophy here is that we are risk-neutral, we are risk-tolerant, we are happy to do the most risky thing if it is the best thing.

0:46:09 Holden: Of course, all else equal, it's better to do things that are more measurable, more likely to work. And so that is just a spectrum that you can pick any point on and that we've picked every point on, and we basically... That is exactly why I think the program officer model is helpful here is because we want these decisions to come from the people who have thought the problem through the most and decided what point in the spectrum is most appropriate. So we sort of think... I sort of think of myself as in a sense the expert in picking areas to work on and people to work on them, and then those people become the experts in those areas and everything about them and whom to fund in them, and then those people are the experts in doing the actual work they do.

0:46:51 Holden: And so yeah, that's the basic philosophy, and we've been all over the map, so like I said, on criminal justice reform we've funded campaigns that have already gotten results; we've also funded alternatives to incarceration that I don't imagine are going to replace prison anytime soon. On animal welfare, we have not invested in cell-cultured meat but we've looked into it. We haven't gone into it mostly because it didn't look very promising, we're still thinking it over. But we have invested in Impossible Foods which is a similar goal using genetically modified plants instead of cell-cultured meat.

0:47:29 Holden: And most of our causes there's some kind of mix, and certainly we're all over the spectrum. There are

certain causes like trying to prevent a massive pandemic or this potential risk for advanced artificial intelligence where kind of the only thing you can do is the really big long-term thing. Yep?

0:47:45 S?: Are there any sort of very messy or controversial issues that you've come across as possible funding opportunities, so examples of that might be eradicating mosquitoes as a disease factor or direct human computer brain interfaces or genetic engineering humans to eliminate diseases or improve intelligence or something like that? Is there any considerations you have when approaching those? Another way to frame this is: Have you encountered maybe trolley problems in your work?

0:48:16 Holden: I don't think I quite followed how the last part connected to the rest. But have we encountered controversial causes? How do we think about controversial causes? I think everything we do is somewhat controversial to someone, and it's just a question of who and how much? Certainly, we are mindful of the fact that if we do things that are controversial, that always incurs a risk that we're going to get bad press, that people that going to form their first judgement of us based on that press, that that will affect how appealing we are as an employer, as a partner.

0:48:51 Holden: So, all else equal, more controversial is not better. It is worse. There are costs to it. And we don't ignore those costs. Generally, we're trying to overall... Our overall goal of the organization is to do the most good. So, to the extent that we are considering controversy, it's as a means toward that goal. But generally, if we find something that we think is really good, we're going to try to find a way to do it. And we're going to try and find a way to just have a good, deliberate, well thought out, communications plan to minimize whatever PR risk there is. Yeah.

0:49:28 S?: How are you working to incorporate constituent voice in potential grantees? And one example would be criminal justice reform disproportionately impacts poor black communities. How are you making sure that you're actually incorporating that into your grants and not exacerbating the problem?

0:49:43 Holden: Sure. How are we incorporating constituent voice, people most affected, into our grants? One answer that is my answer to many things and is a very lazy answer but I think is also a pretty good answer, is that's program officer's job in a lot of ways is again to assess that, to be the expert in their field, to figure out how important that is for their field and how to incorporate it. Because to state the obvious, but I think is a more relevant consideration for criminal justice reform than for farm animal welfare. It's harder to think about how that would work for farm animal welfare. For criminal justice reform, Chloe has I think had a big emphasis on particularly formerly incarcerated persons being leaders, as an important voice, as directly affected people who can speak with authority, who are credible as advocates and who, all else equal, are going to make good leaders of the relevant organizations.

0:50:40 Holden: So, that's been an emphasis of hers in her strategizing and in her communications with grantees and in her choices. And again, that reasoning does not apply in all causes, and a lot of times we're specifically picking out the causes where the people, the persons who are benefiting are the most marginalized, the most disempowered, the least able to advocate for themselves. So, a lot of our causes, it's very hard to make that consideration work. Also, the constituency of people who might 50 years from today be wiped out by a pandemic, some of these are... We picked them because the constituency, it's so hard to even incorporate them in that way. But when they can be incorporated then I think we... Not as an Open Phil policy but as a program officer consideration, to try and incorporate that. Yep.

0:51:28 S?: Regarding the, I believe Open Philanthropy recently made a large donation to OpenAI and some partnership with them.

0:51:34 Holden: Yep.

0:51:37 S?: I've heard a good amount of concerns around OpenAI directly contributing to bad outcomes around AI. My general understanding is that a lot of people at OpenAI are very concerned about the risks, concerned about safety, wanting to do the right things. But the concerns I've had is that what they're working on looks a lot more like general artificial intelligence rather than the narrow applications that you might see at Google.

0:52:05 Holden: Sure.

0:52:06 S?: I assume you don't have a lot of concerns about that given that you guys gave a large grant. But do you have any concerns about them directly contributing to a bad outcome, not just failing to contribute to a positive outcome?

0:52:18 Holden: Sure. So do we have concerns about OpenAI contributing to a bad outcome? Because in some ways

one can argue that they're doing a thing I talked about in terms of being very focused on making artificial intelligence more powerful and not just more safe although the back flip paper I mentioned did come out of OpenAI. So, there is safety work going on there. But yes. Do we have concerns about this? So, OpenAI is definitely one of the cases where we decided to take a big risky swing. Risky in many ways. Very big grant. Definitely things about that situation including what you're talking about that I think from the outside it's a big bet, and the reason that a lot of the reason we did it is that when we look AI and we want to build this field and we want more people working on safety, one of the big problems we saw ourselves as having is that we can fund academics. There's a play book for influencing academia and for helping to build an academic field. But we also feel that a lot of the best people, the best jobs, the most prestige, the most cultural impact is in industry and we just had no idea how to make a difference with industry, because we can't make a grant to Google Brain.

0:53:34 Holden: But we can make a grant to OpenAI. And we did. And so, that was a case of us... It was less about saying, "This organization exactly as it is, we want to just empower them and speed it up," and more a way of saying, "We would like to be part of this organization and we would like the shape of this organization to be partly determined by our input." And we think that is a very good situation for that because we think that OpenAI was started by people who have genuinely really good intentions, who are willing to put aside profit opportunities for those good intentions. And so the basic pieces are in place, and I think they're interested in doing the most they can for safety and want to do the most they can for safety.

0:54:18 Holden: But we've also perceived that in many ways we are maybe more deeply entrenched in the safety community than they are. We've had some disagreements with things they've done in the past which are on our write-up, and so we've seen this partnership as a sort of example sweet spot where there's a group where we really feel that we can be influential in a good and healthy and mutually respectful way, and that's why we did that. The answer to your question is it's not that we have zero concerns, it's not that we think everything is perfect as it is and just needs to be accelerated and sped up. That said, I do have a little bit different picture of the AI situation than I think some of the people have expressed the most concern there.

0:55:04 Holden: I think you will hear a lot of people talk about, who are very into the issues, you will hear people talk about AI as if it's like a race between the two different kinds of research. And I feel that that is pretty misleading way of thinking about it because I think the closer we get to really powerful AIs, the more tractable it's going to be to do the safety research. I think the value of doing safety research today, a lot of the value comes from building a community and having people who are used to thinking about it. I think the value of doing 10 years of pure safety research before you get anywhere in capabilities is quite questionable. And so we do think that a lot of the cultural implications of getting to work with a major leader in industry are more important in our view than that race between safety and capabilities idea, and that's something we're probably going to be writing more about later this year.

0:56:03 Holden: The final thing I'll say about the OpenAI grant, and this is a general thing that's true of just about everything we do is that we do try really hard for our program officers or whoever's leading the work in an area to know the people in the area who know the most, and consult with them a lot. And so in this case we talked extensively with FHI, with MIRI, with other people and groups who we felt were the most knowledgeable and most concerned with the kind of issues you're talking about, and really laid out the pros and cons with them. So this was not something we just went ahead and did. That doesn't make it automatically right, but it's certainly something that we thought about beforehand and talked about with the people we thought would have the most to contribute on that topic. Take a couple more questions, yep.

0:56:45 S?: Do you do any work in promoting democracy or influencing the government around the world, especially in developing countries?

0:56:57 Holden: Yeah, do we do work on promoting democracy especially in developing countries? Not at the moment. So at the moment we feel like we've got our plate pretty full with the things I said, and our main play on the global health and poverty front is GiveWell, GiveWell's top charities, GiveWell's incubation grants, which I haven't thought as much as I'd like to about causes in global poverty, but we've definitely put a fair amount of thought into them and believe that the GiveWell stuff is right up there with what we'd be able to do if we built our own program. So that's not something we focused on today; it could be something we do more of in the future. How many more? I'll take two more.

0:57:35 S?: I asked one already.

0:57:37 Holden: Did you ask one already?

0:57:38 S?: I did.

0:57:38 S?: Not yet.

0:57:39 Holden: Okay, go for it.

0:57:41 S?: So could you tell us about your current thinking or progress on climate engineering?

0:57:48 Holden: Sure. Current thinking and progress on climate engineering. Basically climate change, I think if you took the projections of the IPCC report at face value, climate change would not really look like a global catastrophic risk as we're defining it; it would look really bad, it would look really tragic, it would look definitely worth in my opinion the investment in carbon emissions today to prevent. But the way I've defined global catastrophic risk, it wouldn't look like that. I consider it still to be one because the IPCC report might be wrong because this is a very chaotic system and it's very hard to tell what's going to happen. But our main interest in climate change is in the way out there worst case of climate change, the main place we want to focus, especially because climate change does get a lot more attention from philanthropy than the other causes I've been talking about tonight. So we've decided that the biggest opportunities we see there for us to play a special role have to do with worrying about the very worst case. What's the very worst thing that could happen, worse than any of the models are saying.

0:58:52 Holden: And if that happens, then I think it's quite likely that the world will turn to solar radiation management, which is basically various methods of reflecting back sunlight to actually actively cool the planet. So this could include shooting aerosols into the air to try to actively cool the planet, and this is not something you would want to do for no reason; that's something you'd want to do 'cause you were in trouble. I think that this kind of practice, on one hand it looks like a very cheap and probably fast way of cooling the planet if we needed to. On the other hand it's definitely messing with the ecosystem even more and could bring its own risks, and actually could be the bigger risk than climate change. If things get this bad, you could imagine... There's a lot of disturbing things about solar radiation management, like the fact that in theory any one country could decide to use it and then they would cool the entire planet, which I think is not a great recipe for geopolitical stability.

0:59:53 Holden: And so our interest there has been to say, "If this technology gets used, it could be the thing that sort of saves humanity or it could be the thing that dooms humanity. The better understood it is, if and when that day comes, the better off we are." And so we've been funding this research. Initially we focused on governance research. So, discussions of who should be allowed to do experiments? What is the international governance and coordination mechanism around that? Who should be allowed to deploy this stuff? And more recently we've also been funding the scientific and technical research to say, what is the safest most reliable way to do this? How safe and reliable is it? And our basic thinking is if this day ever comes, we want to be prepared, we want the world to be as well informed as possible because that's a big risk in both directions.

1:00:39 Holden: One more question. Yeah?

1:00:42 S?: I get a second one. Okay.

1:00:43 Holden: Yeah.

1:00:44 S?: Gender. None of the GiveWell recommended charities really look at that as an issue [1:00:50] \_\_\_\_ with that methodology to sorta say that. So when you're looking at global poverty, did you all in your discussions think about gender, and if so, why did you sort of decide, "No, we're not going there?" 'Cause funding streams toward things that even the playing field, particularly globally, outside of Northern Europe, here, wherever, developed countries, is really quite modest.

1:01:15 Holden: Sure. So why haven't we brought more of a gender lens to our grant making? I think the way you posed your question, it might actually be a better question for GiveWell actually. But to answer the question in general, we've tried to pick the causes that look most important, neglected, and tractable. And I think a lot of those causes, they are not necessarily gender-blind. They are not necessarily irrelevant to gender. I think a lot of GiveWell's top charities, if you look at the benefits, they may in many cases be equalizing benefits of a sort, but it hasn't been... I guess the causes we have ended up with don't have gender in the name. And I think that is basically... I don't really have a particular reason

for that, I think there are causes that could have "gender" in the name, that could get onto our priority list later on. But that's just one of many things that is not true about the relatively small set of causes we work on.

1:02:14 S?: And the GiveWell model we've looked at least internal [1:02:17] \_\_\_\_\_ to our research, it's great, it's robust, and looking at gender, it's a very tough thing to measure that way, it almost doesn't work 'cause it's so messy.

1:02:28 Holden: Sure.

1:02:29 S?: I was just thinking I guess because of what you guys do as sort of separate from that methodology of one of the things you've looked at. It's not criminal justice reform but in some ways, it's messy like that. And you have the ability to do that 'cause you have billions of dollars, at least people you work with do.

1:02:50 Holden: It's definitely not, 'cause it's messy, that we're... I would definitely be happy to see if there is anything that our methodology could contribute to that sort of thing. I don't think there's any particular reason that it couldn't, just sort of defining a goal, and trying to work toward it in the way I've described could definitely be something I'd be excited to see someone do.

1:03:12 S?: I have no doubt you guys could do that.

1:03:14 Holden: Okay. [chuckle] Cool. Well, on that note, I will end the formal part, and then I'll hang around for a bit if you want to ask more questions.

[applause]