



Stanford eCorner

Celebrating Failure Fuels Moonshots [Entire Talk]

Astro Teller, X

April 20, 2016

Video URL: <http://ecorner.stanford.edu/videos/4208/Celebrating-Failure-Fuels-Moonshots-Entire-Talk>

Astro Teller, director of Alphabet's moonshot factory, X, describes how smart bets on world-changing innovations are aided by a culture that celebrates only the most audacious projects and rewards teams for showing the courage to find the biggest flaws. He also discusses how innovation can be systematized regardless of business type, resources or role at your company.



Transcript

(applause) - So I wanted to start by telling you something that happened six days ago at X. Which was a little bit unusual, and then maybe we can go back a little bit and unpack how something like that happened. So we have these audacious goals awards. Once a quarter, we get together, and everyone presents how they're doing, and there's a trophy. The trophy hasn't been given out in a year now, 'cause it's really hard to get this trophy and I'll talk a little bit about what it takes to get this trophy. I think people care a lot. I've seen a team cry, and this was a team of patent lawyers by the way, who were weeping they were so happy when they got this trophy. So we had our read back from Q1 that happened just Thursday last week. And very unusually, three different teams potentially merited that trophy. Usually it's nobody, that's why we haven't given it out in a year.

So the star chamber, which was my and my admin Sergio. Went and huddled and decided who actually deserved it, picked the one of the three teams that we thought actually had been the most audacious of the three that had potentially earned it, and then awarded them the trophy at the end of the all hands. And the person came up to accept the award who was sort of the main Lieutenant for that group because the head of the group wasn't there at the time. And he took it, he said thank you, and then he said, but I don't think that we're going good enough, I don't think we've been audacious enough, I'm handing it back to you. It's like we can't give the award away, it's been a year since anybody got it. And this was with many many hundreds of people watching, people have been waiting all year, why would someone hand the award back? So I want to take a step back now, and tell you the thing that excites me the most, this is some of the stuff that I couldn't get into the TED talk which I gave recently Vancouver that's up online now, but that in some way because it was a TED talk, has to skim the surface of the stuff that I'm most excited about and not get into some of the details. I'm a culture engineer. The thing that excites me the most, is not making stratispheric balloons, or self driving cars, or working on contact lenses, or on UAVs that can deliver packages, the thing that excites me the most is trying to systematize innovation. And when I was young, I used to think that systematizing innovation might be some combination of things you could get if you just went to all the business books and you picked out the smartest things from each of the business books. Hire the smartest people, and fail fast and this and that and the other thing.

It's kind of true, it just turns out that if you do that, you don't actually get much innovation. What excites me is, what would it take to actually get a group of people to do the things that it says in those business books that you guys have all read. If this is the set of things they should do, and this is the set of things they do do, and for any of you who've been in business before, you know how big a gulf there is. Wherever you worked, I promise you that gulf existed. There's a reason that there's such a big gap between the things that you want them to do, and the things they actually spend their time doing. It's because this is the lip service that you're giving, but this is this paths of least resistance, emotionally, to doing those things. They don't care what you said they should do, they're going to follow the paths of emotional least resistance. Cultural engineering is the process of trying to get this, the paths of least resistance, to actually line up with the things you want them to do. So I'm gonna give you a few

examples. So and the first one is, let me unpack a little bit more about the audacious goals.

So here's how most companies do something sort of like the audacious goals. Have you guys heard of OKRs? Objectives and key results? So this is the way objectives and key results actually works in a business. You, you report to me, and you're gonna have to held accountable by me. 'Cause I'm the manager. So we're gonna start this weird haggle situation. We're gonna figure out what your OKRs are, you're gonna try to haggle them as low as you can, sand bag sand bag sand bag. 'Cause you know I'm gonna hold you accountable for whatever it is that we decide on. So then I can feel you pulling down, so I'm gonna pull up up up up. I want to haggle haggle haggle, you can do more you can do more, you're sand bagging. And we end up in this place in the middle where you feel like you don't really own that OKR, it's higher than the things that you were saying, and you were making all these arguments about why that's an unreasonable number, or metric for me to hold you to.

It's either the wrong metric or it's too big relative to what you think you can actually accomplish. And I feel bad too, because I feel like it's two thirds or half of what I actually was trying to talk you into. So now we have this thing that neither of us believes in, and this is the OKR. This is a stick, it's a weapon. And my management plan for the entire quarter or year, is I'm going to beat you with the OKR stick. You're not doing it enough, you're not doing it enough. This is 21st century management somehow. It doesn't work very well. This is not how to get people to be innovative. You cannot get them to do the things you really want them to do, especially if your lip service includes things like creativity and failing fast, and being transparent, and a lot of emotionally hard things, while you're beating them with the OKR stick.

Crazy idea, what if instead, you just got to pick what you were gonna do? Let's call it your audacious goal for the quarter. It's your goal, I'm not gonna haggle with you about it, you pick it. You get up in front of all of us, once a quarter, and say I'm gonna try to get this thing done and I know that I'm almost certain not to get it done, but I'm proud of the fact that I'm going to try to do something that sounds so crazy hard, so unlikely. The goal is to have it be something that you can accomplish, about 1/10th of the time. If you're positive you're not gonna accomplish it, that's not very interesting, you're not really gonna try. On the other hand, if you're confident you're gonna do it, it's not audacious, by definition. So you want it to be in that sort of 10% range. And you're going to end up getting held accountable by yourself, 'cause you picked it, and by the whole community, because you want everyone to be proud of you. Now I can be your coach and mentor, instead of having to beat you with the OKR stick. So at X we have audacious goals.

And every quarter, every team gets up and says here is what our audacious goal for the quarter was, here's how we did against it. And here's what we're gonna try to do for the next quarter. And some teams don't do it some quarters, and that's actually okay too. They don't look as audacious when they don't do it, but that's fair game, because you need to be crisp about what you're gonna do if you're gonna try to do it. So now think in that context. That after a year of not handing this out, we offered someone the trophy, and they said no thank you. We're not being audacious enough. Somehow, they've internalized something really deep and meaningful. They were actually one upping me on the things I'm trying to get X to be like. That's when the magic starts happening.

When the flywheel is speeding itself up and so you don't have to keep speeding it up. That's great corporate culture. That's what you want to unlock. In any organization that you're a part of, or that you're overseeing. How did that happen? Let me give you two or three other examples, of the kinds of things that might have lead to this person Grant saying no thank you and handing the trophy back to me. So one of them is fail fast, lean start up. Predotyping, I don't know which one of these things you've heard recently. They're all fine, they're all true. It's that saying it doesn't make it happen. I'm gonna tell you something else that we do.

And afterwards it will sound kind of shocking and yet basic. I encourage you to find anywhere else that has actually done this. When one of our projects that actually has a non trivial number of people, at least a few people full time on it. Ends their project, and they end their projects, I'll tell you a story about that. We bring them upon stage, we have a bunch of Xers here, you've seen this multiple times. We bring them up on stage, and we say this team is ending their project today. They have done more in ending their project, in this quarter, than any of you did to further innovation at X in the last quarter. And then all of you the first time you hear this are gonna feel a little ripped off. Like wait a second, they're failing, they're calling it a day. I'm working my ass off, how come they're up on stage, that seems kind of unfair.

Then I say, and we're giving them bonuses. You say wait I'm not getting a bonus, why are they getting a bonus? They're killing their project, I'm actually succeeding, mine's still going. You know what guys? Take a vacation, and when you come back, the world's your oyster. You'll find some new project to start, or you can pick which project to jump into, depending on which one's going best. At this point all of you are a little mystified if this is your first time through this, and feeling quite ripped off. But social norms are incredibly powerful. People will do horrible horrible things to each other if you set the social norms appropriately, they will also do incredibly innovative, creative, expressive transparent things if you create the right norms for it. By the 10th time that we do this, it's normal. I don't even have to remind people anymore when we stand them up, that they should get a huge round of applause. And that everyone there should be looking up to the people who ended their projects.

But have you ever heard of somebody actually getting rewarded? 'Cause if I tell you to fail fast, are you gonna run out, if you're part of our organization and fail fast just 'cause I said it? No, you're gonna be thinking what happens if I fail fast? Am I gonna get fired? I'm gonna lose all the people reporting to me right? So then I suck and then I'll go tell my friends I was kinda demoted. Am I gonna get my bonus at the end of the year? What happens to my compensation or my opportunity for promotion? Are those things all out the window? This is the difference between the lip service and the actual emotional paths of least resistance. Creating the feeling that failing fast would actually get you what you want instead of getting you the opposite of what you want. So here's a story for how one of these teams got up on stage in the first place. I was sitting with maybe 30 of them, we'd been working for about two months. Their general manager for this project hadn't been working out great, eventually we had to move that person out of the organization. I had become the temporary general manager, I already could kind of smell where things were going but I didn't want to prejudge, I said look, we are gonna decide whether this is going to the moon or whether we're gonna call it a day. You guys need to come to me with that decision, and they were having a hard time doing that, and they were working on kind of some business planning and technical planning, and sort of restructuring their thinking. And eventually, they got me a dock, it was seriously that thick. 30 people in the room, we all sat down.

I think they thought I had pre read the doc, I really hadn't even cracked it. And I said all right who's with me? We're all gonna spend the next eighteen years of our lives if that's what it takes to cause this to change the world. Eighteen hours a day if that's what it takes. We will quit X, all of us, and go do this, who's with me? No one's hand went up. And I just waited, I hadn't looked at the business plan. I just kept waiting. Maybe a good 30 seconds, a minute which was very long when you're all just sitting there and I'm the only one with my hand up. I said sounds like you guys need to talk amongst yourselves a little bit. I went away, they called me about three hours later and said okay come back. And I came back, and they said yeah you're right.

I said we should leave. And they said it's good but it's not great, we don't really think it's that great a business. Awesome, that wasn't so bad. Dee sitting right there was in the room for that, and was part of the shepherd of helping them during the three hours when I wasn't there. Come to that conclusion, but they said it to me. When I say it to them, they've been fired, they've been demoted, their project has been ruined. 'Cause they have no ownership of that moment. But when they say, yeah you're right. They're taking ownership of their own time of their own lives, which believe me is the only really precious thing. When you're pretty good but not great project is over, you're not just gonna go sit in the corner and do nothing afterwards, you're gonna get onto something that's better than that.

Almost by definition if that was a good but not great project. And that they owned that, then they didn't feel like it was being wrest away from them, and then they got stood up in front of X and congratulated, and bonused, and vacationed. And you know what? Most of those people I think like all but one or two of those people are still somewhere in Alphabet, a decent number of them are still in X, there was actually a whole group that sort of recombined around a new project, that actually went fabulously well, and we've now graduated out of X to another part of Alphabet, and I think that they're super happy, and we gave them hugs on their way out and begged them to come back when they're done with that thing and ready for another tour of duty at X. This is not a failure moment, the word failure, and trying to get people to fail, is a bit of a misnomer. First of all, it's missing this path of least resistance issue because people don't wanna do it. Failure, when it's actually just you got a negative result for no reason and it's meaningless, is a bad thing. I'm not pro failure, I'm pro learning. Another thing that we talk about a lot at X, you guys can all probably run this experiment, in your own mind, especially if you've been at work somewhere for at least a few years. Imagine a project that took a non trivial amount of time to do. At least months.

Think about when you were done, you and probably a group of other people, but even just by yourself. When you're done, you lose all the code, all the tools, whatever it is, but you still remember having finished. And what it took to finish. If you had to build it from scratch, whatever it was that's in your mind right now. How long would that really take? To do it again? 10% of the time, as the first time? We have a name for the other 90%, we could call it wasted time but that's not really what it was. You were learning what the right thing to do is. Innovation is that 90%. It's how much you can either compress, the time in that 90%, or compress the cost of that 90%, that's it, that's innovation. And that's why failure matters. That's why we send people out to fail but we don't mean go not be good.

That's not what we mean by suggesting that failure is positive. We mean find incredibly efficient ways to learn. Time efficient or cost efficient ways to learn. And then you have to create these opportunities to feel good about it like some of the ones that I just described to you. Another one that we're trying right now. I think it's working pretty well but, maybe the jury's not totally in on this one yet. Is you guys have heard of post mortems? Maybe some of you have done post mortems. You wait 'til something's dead, and then you look at it being dead and you talk about why it died, and you try to learn something from the fact that it's dead so that you can kill it faster or cause it not to be dead or equivalent the next time. You learn from the death, that's why it's a post mortem after death. We're going pre mortems.

A pre mortem is nothing other than trying to talk about the learning moment of a failure. Before we actually have the failure. We're so eager to learn from our failures, we don't want to wait, 'til the failure happens to learn from it. It introduces a little of one of those time machine movie questions, where if you actually then address it and then the failure doesn't happen, maybe it

never would have happened in the first place. So you don't get a good control experiment for these things. But actually saying to everybody in the organization, let's talk about what's wrong with us. Not in a survey monkey kind of way, but like let's really talk about it. Tell me what you think is the biggest risk for our organization overall, or for project Loon, or for the self driving car project, tell me why we're gonna fail. When we fail, three years from now, what will that be, in your opinion? Write it down, put it up there with your name on it. Which is a little bit scary, because some people can feel thrown under the bus when you actually call out these Achilles heels that you see or that you think might be there.

Then, have a mechanism, which we do so you can just vote these up or down. Which causes the things that more people think are actual risks, even if you didn't write it down. Yeah I agree, I agree, I agree, with these and not with those, no I don't think those are problems. It causes the stuff that's probably the biggest risks, to rise to the surface, and then there's commenting mechanisms so people can actually discuss it. Only if you get thrown under the bus, you say that project Loon is gonna have some problem. And you say what it is, and you work on project Loon, and then people go after you about it, if our culture isn't one that rewards you for doing that, that's the last time you're gonna do that. So making a mechanism like that actually isn't the hard part, I think it's a good thing, it's working for us. But the hard part is relentlessly and repeatedly chasing down those moments where it's not working. He needs a hug if he said something brave on that site. I mean a physical hug, an actual hug.

Or a high five or whatever. And then if he actually gets a hard time from someone for having written that down, what are we all gonna do to defend him, not just because he's right, he's probably wrong, we don't know, there's a lot of smart people on the loon project, I'm sure they've already thought of it. But thank you for saying that. Whether or not you're right. So I'm gonna give you one of more of these things and then I'm gonna open it up for questions so you guys can start thinking about some of this stuff. But let me describe for you a little bit about the process that we have for going through these things, because now that I've given you a flavor for how we're set the expectations at X, it might make a little bit more sense. So we have a team, which is called the rapid eval team. Rapid eval team is supposed to take ideas, from me, from you guys, literally, from the founders at Alphabet, anywhere they can get their hands on an idea, it doesn't matter where the idea is. There's an academic at Berkeley or Stanford or Johns Hopkins, great. Every place is a legitimate place for great ideas to come from.

How can we figure out as fast as possible. That that's a bad idea? That is absolutely and explicitly the question. It sounds like that's not gonna work, we'll just say everything's a bad idea. But if you set the tone the way I've just described it, people are actually interested in coming up with a real reason. Why it's a bad idea. You can't destroy the positivity that comes from saying crazy ideas. But if you say to me, hey, ridiculous idea, do you think we could get the power that's embodied in an avalanche somehow gathered? Maybe that's a way to generate energy. The correct answer, no matter what she said, is that's an awesome idea. She has to feel good about the level of creativity of her idea. If she said something, actually there's 100 companies already doing that, and you purchased something from them yesterday, then maybe that's not an awesome idea, but assuming that it's really outside the box, the correct first answer, the only acceptable first answer is, wow it's beautiful the way your brain works.

Then immediately, that's so great, how are we gonna figure out that that's a bad idea? That that's not gonna work. So she just got a little check mark, with myself, with her peers, for having said something that was really interesting, that was innovative, that was different, than what we were thinking before. And immediately she now also gets to get another check mark if she can show the intellectual rigor for why it's a bad idea. Well okay I guess we could try to generate avalanches, how much is in an avalanche is good, it's not great. Okay well maybe we could move the thing that's gonna turn all that potential energy and kinetic energy into watts of stored energy. Maybe we'll move it around so we can get the avalanches as they fall, no that's not really gonna work. It won't take us but five minutes to sort out that there's probably no practical way to do that, good, awesome, we've figured out rigorously, not just in our gut, that that's not gonna work, and we can move on. Because the rate limiting step to innovation is not finding smart people. You're all plenty smart enough. It is not being creative.

How many people here in this room think that you're highly creative? Good. The other half of you are wrong. You're all highly creative. How many of you think we were creative when you were six? Who wasn't creative when they were six years old? You don't have a six year old if you think that you weren't creative when you were six. We just get it beaten out of us by society, I promise you, you are creative when you were six years old. We all were, we've just forgotten how because the context isn't inspiring us, isn't allowing us, it's literally blocking us. But that's not the problem. The problem is how to get a huge number of ideas on the table, and then weed through them effectively which is not about process, it's about creating an environment where people feel like they can be rewarded in emotional ways and financial ways, for doing that. A tiny fraction of these ideas then pass through to our sort of second stage booster rocket, which we call the foundry. In the first stage, most of the de-risking that we do is on the technical front.

Building prototypes, verifying that it's not some isomorphism to like a perpetual motion machine, you'd be surprised, probably one in 100 of the ideas we get, it literally is an isomorphism to a perpetual motion machine. Once something gets to foundry, maybe 20 30% of the work is still very technical, but a lot more of the work then gets applied to what is the ecosystem

like? And the regulatory environment. How much would we have to invest versus how sizable a business would this be. How much good would this really do for the world? If we didn't do this, would the world end up with that benefit anyway for some other reason? Or not? Working through all those reasons, again for the purpose of killing the project. Even in foundry, which is supposed to only receive things, that have been heavily weeded. The goal is to have more than half of those projects be killed. When you hit more than half, you're clearly in a mode where the people in foundry understand, even though they can be very passionate about the projects they're working on. That there's less than a 50% chance of doing it, so they can take pride in ending the projects for the right reasons. Eventually, for the self driving cars I'm pretty sure cars are gonna drive themselves since cars are already driving themselves. For project Loon, we have a lot of balloons up in the air, they're already doing LTE to the ground, people are actually receiving phone calls so we know it's possible, kind of the ship has sailed, pun intended a little bit on some of that stuff.

But for a long time, the pressure is not how can we make this work, it's how can we discover as fast as possible, this is not gonna work, so that we can get on to doing something else. I'm happy to keep going, but do you guys have questions? I want to make sure that I shape some of this for the stuff that you guys are thinking about doing. All right cool points for the first question. Yeah. - [Voiceover] How would you apply this to big projects? Namely something where you want to put a man on the moon. Or do a Manhattan project. Or have you thought about, say Einstein's comment that it's 90% perspiration, rather than 10% inspiration. In other words, the devil's in the details, the hard part is doing it, as opposed to thinking of it. Ideas are a dime a dozen. - Right well actually I agree, I just gave you an estimate which was also 90%.

Believe me we do a ton of perspiring. I'm describing that 90%, the learning part, as the perspiring. Because that 90% you guys all know from whatever it was that was the last hard project you did, you didn't just sit there and navel gaze for that 90% of the time, you thought you had the right answer. And you perspired and you did stuff and it didn't work out. And you're like oh man that didn't work out. And then you have to learn and iterate. That process, which is the perspiring, that's what I'm talking about. What I'm describing, is exactly the process of saying whether you're making a rocket to go to the moon, or you're making a car that drives itself, what might cause this to fail? Let's set up hypothesis, your questions. Let me give another example. I've run this experiment many times with different groups at X.

Make me a list of the top 10 things that your team needs to do. No problem Astro. And we get this list, and it's invariably ordered, by the things they think will be most important, will be most well received by the market, whatever that is. Thanks. All right now humor me. Reorder that list, instead of most important at the top, make it the thing at the top is the one that would cause us to learn the most. And the one at the bottom is the one that will cause us to learn the least. They reorder the list and is invariably a different ordering of the list. Okay. Why don't you just do the first two things on the list? Just humor me, just do the first two things on that new list.

And then we do them, and then I say, make me a top 10 list of the things you think are most important. Forget the learning stuff. And they do. And that new list is unrecognizably different from the previous list of 10 things that they made. And then I say that's why we focus on learning. Because if we hadn't done that, you would've done all 10 things on that list, and somewhere in that list, you would have run across the things that actually caused you to learn, such that the rest of the things on the list turned out to be irrelevant. Why wouldn't you try to bubble those two things to the top of the list. Always, this is not about being smart. This is not about being smart. It's not about being creative.

It's about choosing to spend your time on that. And I cannot over emphasize, you cannot just tell people to do what I just told them to do. Feel free to take the idea with it. You will discover that it doesn't work. Not because it's not true, but because people won't do it. It will be maddening to you, I promise you as a young entrepreneur you will discover that people don't do it, and it will make your head explode. But they won't do it. I'm speaking from personal experience. The reason they won't do it, is because put yourself in their shoes. If you don't have someone who's ranting the way I'm ranting right now, if you do the things on that first list, the important stuff, you get the important stuff done.

After a couple months, you bring it to your boss, to the CEO of your business, and you say look what I built. You got important stuff done. Good job man. Hey look what I built, if you follow that second thing. I just proved that we're totally wasting our time and we're gonna have to start over. You really think that they're gonna get a thank you. 'Cause they don't think they're gonna get a thank you. And if they don't think they're gonna get a thank you, they're not going to do it. That's it, that's the hard part. It's not the reordering of the list, it's creating the culture where people are willing to listen to you about the reordering of the list.

Other questions. Yeah. - [Voiceover] Thinking of this from my post mortem point of view, what do you think are the biggest challenges or causes of failures for Skunkwork organization like X to operate within a bigger company such as Google? - One of the challenges for doing something, yeah sure I will repeat it. The question is, what are some potential failure modes for doing something like X inside a large entity like Alphabet. And I'm not sure that this is a complete list, but if you are coupled too tightly to an organization. If you're busy trying to solve their problems for them, then their needs can whipsaw you in various ways. That can be productive in the good cases, where you actually solve one of their problems then they'll really appreciate it.

And you have those sort of ready market for your thing that you've made, 'cause they already needed it. And you don't have to think about what's hard and important, you just have to think about whether you can make the thing that they need. But in the end, businesses tend to be somewhat sigmoid, right.

They're flat for a while, they go up really steeply, and then they taper off. In the long run. And they way businesses continue to grow in the longer term is by having more than one problem to have. So X's mission is to find some new problems for us, more generally to have. For the larger entity to have. And if that's our goal, to find new problems for us to have, we can't be overly constrained by what they think their problems are today, so it's a failure mode to be too tightly coupled. I also think it's a time horizon issue. Everything that I just told you, I believe very deeply. But it's probably not the right thing to do, if you have to declare victory and cash in in three months. Over sufficiently short time horizons, empowering people and trusting people is probably a waste of time.

Over a very short time horizons, you should probably just find the rock stars, and use fear and greed to get them to get stuff done. I think that's probably, over sufficiently short time horizons, I think that's probably the right thing to do. The opposite is also true, if you have decades, you cannot over trust and over empower people. Even if somebody spends a half a year learning and growing and they waste some the entity's money in the process of doing that, if you really have many decades before you have to declare victory, you'll get paid back many times over, because of that growth, because of how empowered they feel. Now the truth is time horizons are usually somewhere in between weeks and many decades, but if you set the time horizon wrong. I think that can also potentially be a failure mode. X's time horizon I would describe like this. We're looking for things that can really have a huge meaningful impact on the world within about five to 10 years. But, if somebody were to bring me, and we do this experiment frequently. When we have two things, this one maybe could change the world in four years, but we won't know for three and a half if it's any good.

This one could take 10 years to change the world. But we'll learn something every six months about whether we're on the right path, forget that one. The difference between five or 10 years, we're kidding ourselves that we can even tell right now, whether it's gonna take five or 10. I want the one where we can learn along the way, not where it's this one shot thing where we find out after three and a half years that we've been wasting our time or not. So I would say that's another potential failure mode is ending up where you think you're a skunkworks, but you're not a learning organization, you're a take big risks blindly organization. Which is probably not how I would suggest doing it. Yeah. - [Voiceover] What might be the most impressive proposal or idea that's come your way? - Impressive proposal. Oh yeah so the question was what's the most impressive proposal that's come my way. You know, this might be an unsatisfying answer.

But I don't judge it that way and I don't I hope nobody at X does either. I think you might be surprised how much it feels collegial at X. Because as soon as you allow too much one upmanship to enter the process, dangerous things happen. People start trying to get that award, and we don't want them to want that award. I gave you this example before about the avalanche. I like that as an example, because it's different, not because she was right, but because she was way outside the box. Someone else, it wasn't her, but someone else asked a somewhat different question which was, could you put a ring of copper around the north pole or the south I guess, and use the Earth's magnetic flux right, the molten core of the Earth kind of sloshes back and forth a little bit, and you get sort of moving B field, which would generate a current in this huge piece of copper that you put around Santa Claus's house, and then you could pipe that current back down here. It turns out not to work, but I mean it works, it's just not at all cost effective. But it actually would generate electricity. AC obviously.

But very slow wave AC. I love those questions because they make me believe that we haven't even scratched the surface of our collective creativity as a species yet. And it's important because it shapes something about X. Part of what I've just described about killing these projects as fast as we can, has a hidden assumption in it, which I believe very strongly, but you have to believe this in order for what I've just described to be right. The cost of a false positive. At X, at any innovation factory, is huge relative to the cost of a false negative. A false positive is where we think we have a great idea and it's not a great idea. That's a false positive. Well if we think it's a great idea when it's really not, we're spending money, we're spending money as all this management time the team keeps growing, and we're wasting all of our money because at least in this thought experiment, it's a bad idea. What's the cost of a false negative? If we say that something, which is actually a great idea, we look at it and we say no that's not a great idea.

That's only a cost to us if there's a really limited supply of great ideas. Of really big problems with really innovative solutions. Sadly, humanity is no danger of running out of huge problems to work on, and I really don't believe that we're in any danger of running out of really interesting incredibly different potential solutions. And because both of those things are so rich and such great supply, that false negatives cost us almost nothing relative to the false positives, better to lean fairly heavily towards the false negatives rather than the false positives. Yeah. - [Voiceover] (mumbles) What do you do with the things (unintelligible talking). - So the question is in these pre mortems, when a bunch of people vote up something, there's now 50 or 100 people who've said yeah I agree, that might be a big risk for X. Or for McConney, our airborne wind turbine project. First of all, we take it seriously, like we look at it more. Number two, it depends because sometimes they have very different

characteristics.

If it was, we're gonna run out of space, then that's a facilities issue, and that's not like kind of a panic, like we need to tell the facilities person that, like McConney might not even need to be involved in that particular case. Or if it's more of a whistle blowy thing, where it's I think we're going so fast that we might be cutting corners on safety. Okay you can't just send that over to McConney as a whole and expect the right results. You might want to cherry pick the person there who's responsible for safety. And say what do you think about this? What kind of pressure are you getting? Give me examples of the last three times something bad happened, like how bad was it? So that one might be a bad one to just send to in this case, Ford is the general manager for the McConney project. If it's a safety thing, maybe just sending it to him isn't the right thing to do but for some, as long as it's not for it's a bad person, and I think he needs to go, then we should just give it the fort and say hey, this is a real concern, what do you think? Often, not always, they actually will then respond on this pre mortem thing in the conversation. And frequently they will ask for more feedback. So if the concern is, we're doing the wrong thing. Okay thank I appreciate your concern, that does not really tell us very much about how we should address that, can you be more specific about what you think we're going wrong? We're going too slow. That one happens a non trivial amount.

Okay that might well be true but that doesn't really tell us what to do about that, can you be more specific before we understand how to route it to people. Yeah. - [Voiceover] So you mentioned two different questions you ask when you get these ideas. How is this gonna push the boundaries, but also how is the market gonna receive this. So do you ever get ideas that maybe don't score evenly on those two questions? And maybe do you have something like a device that can harness energy around the sun, and in fact 10 years you could learn a lot about that device and you could really push the boundaries and technology, but in five to 10 years that wouldn't be ready for the market. Is that still something you would explore at Google X? - So the question is what happens if something looks really incredibly promising on the technology front, but weak, on only very delayed, maybe it's great but 20 or 30 years from now. On the business front. Or vice versa, it looks like a great business, but it's kinda weak on the technology front. In both cases, again because of our bias about the difference between the false positives and false negatives, we mostly say no. We're not perfect, reasonable people can disagree, often there are some people who think oh there's nothing hard on the technology front here, let's not work on this.

Then someone else says, that's because you're not a mechanical engineer, you don't understand that this is like totally rocket science to make this thing happen that's just been proposed. So we do have some of those debates, but once they net out, usually if they don't have both of those things in healthy doses, we won't do it, and that means we throw out 99% of the stuff that even that we talk about seriously. But it's still cheap. Because 100 hours on every single one of these things throwing out 99% of them is still rounding error on the amount of total energy that we spend at X, on the stuff that becomes mature. So if something is 20 years out from actually being something that we can bring to the world, have real impact on the world, I wonder whether that's the best way for us to help humanity, relative to finding something that's only seven or eight years out, working on that, waiting and then 10 years from now, let's get back to that idea of yours, when there's only 10 years 'til it's gonna change the world. Yeah. - [Voiceover] How do you organize your knowledge so that most people can benefit from it? - So the question is how do you organize the knowledge so that most people can benefit from it, there are two aspects to that is, when we come up with an idea internally or someone else gives us the idea and then we go through the 100 hours worth of work, and we decide it's not the right idea, where do we write that down internally so we don't just keep reinventing the square wheel over and over again. No it's a square wheel that's no good, when we've in fact sorted out like 10 times that square wheels aren't any good. We're okay but not great at that. We have repositories internally for doing that, so it's not that we don't write it down, but how it gets shared could be better, the thing that kind of saves us from that at least internally is that we have a relatively focused group, the rapid eval group.

So they all know where the repository is. So people in other parts of the organization probably don't know that it's there, but the people who are really responsible for figuring out in those early days, is this a round wheel or a square wheel? They've written down their own failures and everyone else's failures, sort of these post mortems about why it actually didn't work out. Separate issue is when we discover that something isn't a good idea, how do we tell the world about that? It's complicated in various ways for intellectual property reasons, there's too many to talk about, some of them, like the one you just conceptually brought up, we might not really have killed, we might just be pausing for 10 years, but I just mentioned to in my TED talk, recently one was about these lighter than air variable buoyancy cargo ships. And so we were just sharing with the world some of the basic reasons why we decided not to do that. Someone else might decide to do it, another one that I mentioned was vertical farming, and why we got stuck on vertical farming, why we couldn't get past a certain point on vertical farming. And we've had a lot of people somewhat bristly come back and say well but I still believe in vertical farming. Hey that's cool, we're just telling you why we didn't. The reason we didn't do it, is because we couldn't figure out how to get staple crops. Like grains or rice to grow in vertical farming. If you want to grow like micrograins, go crazy, vertical farming is awesome.

That might be a nice business and it'll seem cool, you're not really gonna change the world or feed billions of people, making micrograins and some 30 story farm, that's not potatoes, rice, corn, wheat. Those things are the kinds of things that if you can't do any of them, it's a little less plausible that vertical farming would change the world. So we just described that to the

world as a sort of, and in fact in my TED talk, I told people explicitly, if you know how to do this, come tell us 'cause this is what we got stalled on. If you have the missing piece for us, we'd love to resuscitate that project. So we do do that, I think it's not out yet, but we just had a big project, that wasn't huge but it was really near and dear to our heart. There was a good number of people working on it for kind of a while, inside X, who really wanted this one to work out. We had some working prototypes, it just felt so good. And we've killed it recently, and we're publishing a paper, anyone have the date on roughly when that's coming out? Not this year? So all right. It's like Nature or something like that. So as soon as we get it through that process, it'll be out to the world, so we are serious about sharing.

Yeah. - [Voiceover] (mumbles) and traditional business. How is some of these principles be (mumbles) For example (unintelligible talking) - The question is how much is what I'm describing transplantable or not transplantable into an organization which isn't trying to be a moon shot factory. Which might be a smaller business, it might be a less technology focused business, or whatever that is. I don't think there's anything about what I said. We're focused on working on things that are big hard technology problems, but that's a detail, that's specific to us. There's nothing about what I said that you could not do except in the presence of big technology problems. The technology is just happens to be an instantiation of it for us. But here's an example. Which is another sort of example of the culture.

In the sort of mid days of glass, I was worried that the team was going to get less creative as we moved prematurely into trying to ramp up manufacturing. So you have to take a bunch of people who are hyper creative, and get them over the wall culturally into this sort of high execution mode, which is very uncomfortable for a lot of them, and you're bringing in new people to help with that process, and this often kills the very thing that you care about in these early organizations, which is the sort of rapid prototyping, and the creativity. So I started a process which we called the get weirder award, And the get weirder award was given out once every two weeks, and in order to make sure it was clear to people what we wanted to reward the reward was given out for the most creative proposed experiment, and we gave out the award before the experiment was run, so that there was absolutely no question that it was not about outcomes, but outcomes you cannot control. As a scientist, you only care about the quality of the experiment. How thoughtfully did you design the experiment? How cheaply could you create the experimental apparatus for the experiment? And by experiment, I did not necessarily mean technology. This is actually answering your question. People from the customer service group won the get weirder award. And from the sales group, and from the marketing group. Technology people did, designers did, but you can ask a super creative question. So one of the examples that happened was we had a one day sale for Glass.

And we couldn't possibly staff up enough people to man the phones and answer pre sales questions. So the customer service group said why don't we ask all of our existing customers, the explorer, for Glass, if they would be willing to do pre sales question answering. It doesn't matter whether it works. That's not the point, that was an awesome question, right? And has nothing to do with technology. And your organization will be successful to the extent that you get people doing that. Asking those kind of questions, and they will only ask those kind of questions, and run those kinds of experiments if they actually think they can get promoted for doing that, and having it not work. Which by the way will almost exclusively not work, it happened to work in that case. It worked really well, but most of these experiments, the very ones that are weird, won't work out. If they all worked out, they wouldn't be that weird. People would already be doing it.

So you have to reward them beforehand, you have to reward them, even if I might even say particularly when it doesn't work out, or they just won't do it. Yeah. - [Voiceover] I'm wondering about how the structure could be implemented in groups where you don't have multiple outcomes like promotion or like (mumbles) or whatever. How can you actually culture the engineer situations where there are counter productive social norms. - So the question is when you don't have some of the classic structures for awarding people, like money or a promotions, how can you do some of this culture engineering? People don't even respond the strongest to money. There's a ton of evidence out there that when you give people a raise in their salary, it has this ridiculously short life, they feel good for like a week. Maybe, and then they totally forget about it. What people want is recognition, they want you to say thank you, they want to believe that their manager, that their peers think they're cool, think they're interesting. Think that they'd be worth having on the next project. Think about yourself, isn't that what motivates you the most in the end? Is being around people you trust, and having them think that you're special to them.

So if you can create any environment in which people can say to each other, hey I really appreciate that, that makes me think you're special. And there's lots of ways, you just give out cool points. Like make little pieces of paper that say one cool point, and then hand them out, or you don't like that, do something else, it doesn't matter what it is, but you don't need cash and you don't need promotions to do what I'm describing. Yeah. - [Voiceover] You spent a lot of time talking about culture engineering. Which is really fascinating but, it seems like culture engineering requires a lot of power for you to actually be able to (mumbles) these big changes. Most of us going into our jobs, we're not gonna have the ability to do all the things you're describing, we're towards the bottom so, I ended leaving my last job, that I worked at for several years because it was really old school and not very interesting. So what do we do to try to implement this culture from the bottom and not the top? - So the question is, if you're not the CEO, how do you implement, or change the culture? And that's a great question. Number one, good for you for leaving a boring place. That is in itself an expression of power.

By the way. That is a powerful thing you can all do, and you are disempowering yourself if you don't do it, and I encourage people to leave places, including X sometimes. If this is not the right place for you, don't humor me, tell me this is not the right place for you. I'll help you find another place. In some other part of Alphabet or somewhere else in the world. I definitely don't want you here, if you don't want to be here. If you think that we're idiots just tell me, I'm actually interested. If you think we're an idiot, and by the way, people tell me all the time that I'm an idiot which is a little bit hard to hear, but in and amongst the you're an idiots, I also hear lots of things that are wrong with X, and we fix it. Like X is the worst moon shot factory in the whole world, except for all the other ones, and the reason is it is because the people believe that it's same to say what's wrong with it, and then we surface all of the bad stuff and very slowly, we're like taking care of it. So I just had a meeting today, I'm not the proud general manager of a different group.

At X. That has some leader challenges. And I was sitting with someone today I'm doing a bunch of office hours, and someone came to me and said hey how are you doing? I don't really know what we should talk about. And I said well what's broken? Oh there's some stuff that's broken. Well give me an example. Oh Bob over there is working on something that we could buy for \$1000. Have you said something to Bob? No. (laughter) Come on and Bob literally happened to be coming to my office hours like 30 minutes later. And I said I'm gonna tell Bob that you think that. I'm not throwing you under the bus, but shame on you for not having gone and talked to Bob.

You should run and tell Bob before I tell Bob. Don't you want to be part of something where we're all like passionate, and we have enough of a sense of urgency that we can't tolerate that kind of waste? That's not just me having power, I'm empowering him, yes, but I'm empowering him to be part of the change that he wants to see. And even if I hadn't done that, if he had come to me and said. You know what Astro it's been a little crazy here, I've been along for the ride and I kinda want to believe, that this particular team within X. I don't know I have my doubts, I think our culture's gone the wrong way, I'm this close to leaving. If you have my back, I will make a fuss, I will go rally us. Me, I'm at the bottom of the totem pole. But I still Bob or Sally or whoever, am willing to be one of your embasaries, to change this and make it great, but you need to have my back Astro. Ask for that, and if they have your back, go do it. And if they don't, quit again.

Go find someone who will have your back. Or just be the CEO, I think half of you are sitting in this room 'cause you planned to just be the CEO so you don't have to ask permission. Which is its own kind of power. - [Voiceover] I'm sure you'll all agree, this was totally fascinating. Please join me in thanking, - Thank you. (applause) Thank you very much.