



## Stanford eCorner

### Forecasting the Future of Technology [Entire Talk]

**Steve Jurvetson, *Draper Fisher Jurvetson*; Astro Teller, *X*;  
Christina Smolke, *Stanford University School of Medicine***

**October 07, 2015**

**Video URL:** <http://ecorner.stanford.edu/videos/3568/Forecasting-the-Future-of-Technology-Entire-Talk>

A special panel of highly scientific minds discusses what the future holds for tech innovation, education and entrepreneurship. Panelists include Google's "captain of moonshots," Astro Teller, Stanford bioengineer Christina Smolke, an associate professor at the university's medical school, and DFJ General Partner Steve Jurvetson. Persis Drell, dean of the Stanford School of Engineering, moderates the discussion, with introductions by Stanford Professor Kathleen Eisenhardt.



#### Transcript

So I get to have the fun of asking the questions to this really great panel with three, I'd like to point out, really different points of view. Steve's venture capital funding, cool new ideas. Astro's running probably one of the most inventive labs, but I think it is fair to say managing a lot these days. Maybe still there inventing. And Christina, she's in the lab making the future happen. Though the future is the theme, and so I want to start out by asking our panelists to look back. Now, because "Back to the Future" came out 30 years ago, I was going to ask them to look back 30 years. And then I realized that all of you in the audience weren't even in existence 30 years ago and that wouldn't mean anything. So we're just going to ask you to look back maybe 10 years. And maybe I'll start with you, Steve.

Name one thing that didn't happen that you were sure was going to happen 10 years ago? Oh, wow. I knew you were going to cold call me, but I still didn't prepare a thought. There's a bunch. 10 years is tough, because there have been so many mistakes that I'm trying to timestamp them. So during the peak-- Just one. OK. I thought Amazon was going to be an anachronism and totally made obsolete by distributed, federated, personalized systems. I thought Netflix would fail because clearly video was going to come over the internet, not these DVDs they were mailing at the time. So these industry transitions, like a company that starts in one area but expands or pivots or grows, I just didn't see those coming. And I think I was maybe at the tail end of being wrong about nanotechnology and the rate at which it would impact the world and how it would, and whether the term meant anything to anybody but me and a few geeks like me.

So that was a whole section of mistakes leading up to 2005 around that sector. So Astro, you get the flip side. What did you predict that is happening? What did you get right? Oh, he gets the same one right. I'm with Steve. I think it's more interesting-- I mean, I'll answer your question. I'll pick something. But I want to start with one or two things that I totally got wrong. Fine. The one that sticks in my mind-- I think of this all the time, because I have this mantra at X about how almost all of my ideas are wrong, how all of your ideas are wrong, all of our ideas are wrong almost all the time. It's how we claw through all the piles of ideas that we have to find those gems.

That's the real trick. And the thing that I use to remind myself all the time how wrong I can be is I remember seeing the first phone about 10 years ago that had a camera on it. And I just thought, people are not going to take pictures with their phones. And even if they did, there's no way that they're going to be like, obsessed, sending pictures to their friends. That's like the dumbest thing I've ever heard. And I was really sure that that was not going to make it. I just remember very viscerally being certain about that. So it's a good humility moment for me. So was there something that sticks in your mind that you got right,

that you could see the future 10 years ago? I think that-- so my background is artificial intelligence, and 10 years ago was a so-called AI winter. And we're out of that winter now.

And that was one of the things I was incredibly confident about, because AI has been on this very predictable sort of exponential increase in its abilities. And that people were saying AI is dead or that it's never going to actually accomplish anything really never changed that trajectory. So that was one of the things where I felt like, yeah, we'll see. And in fact, AI has continued to make really steady progress in a bunch of areas, like machine translation, and speech to text, and computer vision, and robotics, and a bunch of the areas in which it's just continued to churn out benefits over the years. So I feel like I was right about that one. OK. And Christina-- something that you couldn't have even imagined, didn't think about, didn't think it was going to be right or wrong, but it's happened, and it's changing the world. Biggest surprise. Well, I guess I would say-- let me just say two things. So something that I didn't even imagine people would want to work on beyond sort of fantasy and movies that is not necessarily here yet, but people are seriously working on, is de-extinction.

And again, it's not that we've achieved that yet, but people are inching closer to that. And people can kind of see this happening with different creatures and animals. And you know, it just wouldn't have even entered my mind that that would be something that we'd want to spend our time doing. The other thing, though, that is now here that I wouldn't have predicted or thought about is the idea of cloud laboratories within biology. So really now thinking about outsourcing our experiments and the way that we do experiments and examine things to companies where that can be more systematized would not have been something that I would have even thought was possible, or that there was even really a market for. So now let's turn to the future. Steve, what's the coolest thing in the pipeline? What time frame? Like, right now cool, or-- You call it. Oh God, it's all over the place. So the high-level observation would be that when I started in this business of investing in startups 20 years ago, it was all in a few industries of software, semiconductors and biotech. And we didn't do biotech at DFJ, so it was software, semiconductors, and variance computing, photonics.

It all felt the same. And today it's a whole variety of industries that are being transformed because of deep learning, machine learning techniques. So everything from rockets and drones and robots and satellites that used to be really bad industrial businesses, they're now becoming software-centric businesses. So in every one of those sectors, it's kind of a recurring theme of how they're doing it, using off-the-shelf hardware from cellphones repurposed in new and interesting ways. And that's exciting. And then closer to Christina's area of focus, the whole revolution of learning going on in life sciences and how that percolates back to IT in general. Like how do we build complex systems? How will we build these AIs of the future I think will be informed by the way biology is done, both as an existence proof and metaphorical inspiration for how do you iterate on an algorithm and compound intelligence over time, or complexity over time. I think those are the metaphors of the future for how we do engineering. I'm gonna interrupt you for one second. Can people hear OK? Raise your hand in the back if it's OK.

Good question. It's OK. Hand is OK. Good, that's rare for me. Good. Thanks. No, 'cause I think this morning it was tough, so thanks for checking. And so I think the field that excites me the most is biology. The technique is machine learning and deep learning across all fields. And then that leads to all kinds of things, like satellites that will observe the Earth every day, and broadband everywhere on the planet via satellite, and building new life forms that do interesting things, building organs for transplant that are not gonna trigger the immune system.

There's just probably 18,000 different student projects that would have seemed like science fiction even five years ago. So Astro, can I ask you that question? You just did. I mean, obviously there are specific things that I can't answer, but one of the things that excites me which is related to what Steve said is, when I look at the commonality across a lot of the things that we're doing at X, I notice that a lot of the projects are both non-traditional in being robotics but in fact actually are robots. So we're trying to get the other 4 and 1/2 billion people who don't have a good internet connection today an internet connection via a network of stratospheric balloons. But in fact, those balloons have to do a lot of sensing. They have a lot of actuation. They can rise and fall. They're trying to pick the winds that they grab and cause themselves to sail to keep themselves in the useful places. So if you have sensing, you have computation, and you have actuation, in the very general sense, you're a robot. Our airborne wind turbines are actually this shift from a huge amount of steel to mechanically solve a problem to essentially a very small tether and a super hard control system problem, which is essentially bringing robotics much more centrally to this issue of generating energy.

Our cars-- self-driving cars are kind of the poster child for nontraditional robotics. Our self-flying vehicles for package delivery are very centrally sort of robots. So in the sense of robots as walking mechanical people, I don't know how fast that's really going to change the world. But I see a lot of these issues about bringing techniques from machine learning, from mechanical engineering, from ways of embedding intelligence into products as an alternative to making them mechanically safer or mechanically cheaper. I see that as a very general thing that we're going to be able to do over and over again to solve problems in the world. That's cool. Given this revolution that's going on, all the changes-- Christina, you work in a lab. You do your research. How do you choose from that whole menu of really exciting things what to focus on and what you do? So I think

the first thing for any of us is to really ask yourself, ourselves, what are we passionate about and what do we want to be doing? Because when you think about what problems you want to work on and what type of world you want to live in, there are many different routes to take. But you want to choose something that you're going to want to put your whole self and your whole heart, body and mind into.

So that's the first thing, I think, that really drives that decision. The second thing, when I think about problems that I want to work on is-- and especially operating within an academic environment, I mean we-- or at least I think our role is to not just think two to five years into the future. Our role is to think 10, 20 years into the future. We should be doing things that are going to build the foundational technologies and advances that then eventually will move on into industry. So I think we want to look 10, 15, 20 years into the future and ask ourselves, what are the big problems we see, what type of world do we want to live in, and what is either the foundational research or the applied research that we want to focus on that we believe is going to make and effect transformative changes in the world around us? And then you pick what problems you want to work on, and you dig in. And the great thing about working in an academic environment is you have so many people around you to work with that are going to bring their creativity and spirit into it and take things in entirely-- in directions that you couldn't even imagine. And you have to be able to give them the freedom to do that. OK. Steve, how do think policy it's going to play into the future innovations we're talking about here? Government policy? Yeah. Gosh.

It's kind of like the Hippocratic Oath. So first do no harm, I guess, comes to mind. Does that mean you don't work on things-- No, almost never. OK, let's just clarify that. No, no, no. So our companies interact with policy a lot, but we rarely if ever bet on policy swinging our way, with possibly the exception of autonomous vehicles, where we don't how it'll happen, but it just seems so compellingly obvious that all vehicles will be autonomous and electric in the future, and we'll marvel at the present day that we let people slaughter each other on roads. I'm one of those high-risk category kind of people myself, so that and planes. So we know that's the future, but in general let's say in the energy, clean-tech investments that we've made, we never count on policy making the investment work. And so we generally are pretty hands-off when it comes to that. Now government can play a big role, obviously, in sponsoring projects, whether it's NASA helping SpaceX get its start, or a big loan to Tesla, obviously, at a critical juncture.

There have been really important moments. But rarely is it one of sort of company picking. That's the rare model. Usually it's NSF grants and government grants to universities that are the bread and butter of how we see government helping the tech world. And in policy, we just kind of hope they don't make any more. You know that's wishful thinking. I know. I just kind of try to keep my head low and hope they don't bat it down. We've had sort of the opposite experience at X, which is we're trying to pick things that we want to fix about the world, whether it's cars that drive themselves, or the contact lenses that we built for monitoring people's bodies, putting computer and sensors inside contact lenses. So whether it's the FAA or the FCC or the FDA or the Fish and Wildlife, we've had to talk to all of them.

Fish and Wildlife? Well, for our airborne wind turbines. Oh, yes. There's no part of the government-- the National Highway Traffic Safety Administration-- that we haven't had not just to visit once or twice, but pretty much go live with them. But our experience has been that if you wait until you're done making something, and then you go stand in front of them and you say, is this OK, and you don't really mean is this OK, you mean this will be complete meltdown in my life and the life of lots of other people if you say anything other than yes. Because you're not going to get a straight yes. Then you don't get the straight yes, and then you come away feeling like these people are against innovation. Which I don't think is fair to them, because they're trying-- in an ideal world, regulators would have a membrane that prevents all the bad snake oil from getting through but lets all of the good innovation go through really fast. But of course, you can't tell ahead of time what's the snake oil and what's the good stuff. I think they do a pretty decent job. Our experience has been, if you go to them several years before you're done and you say, this is what we're working on.

What do you think? Here's the goal that we have. This is actually the problem that we're trying to solve. And we're not hung up about our solution, because it's not done. We're actually really hung up about solving this problem in the world. And the regulators universally are excited about the fact that we're trying to solve problems for people. And then they can engage with us in a creative process of saying, great, how do we help you to get these things done so it can be safe for everybody, and you still don't feel like you're being stymied by us? So I think it depends how you go about it. I agree, we would never bet on the regulations changing. But then again, we almost always bet that regulators can work with us to change in ways that are going to be necessary in order for the world to be a better place. I love the optimism of that answer. Having spent quite a few years going back and forth with Washington, I can't quite be that optimistic.

But I'm going to subscribe. I'm going to subscribe, because this is the right way to do it. Depends how you ask. Absolutely. So I'm gonna switch the topic now, because this room is filled with students. And I'm anticipating that you three might have different views on this, but I'd kind of like your thoughts on how do we educate the engineer of the 21st century. What's the view of a broad education versus, should we just become Stanford Tech? I mean, 30%, 36% of our students are graduating in engineering. Let's just-- why not all of them? And what do you think is important about what one learns in college and carries

forward from it? And so maybe I'm going to give it over to Christina for the first try on that one, as the educator who's actually in the classroom teaching, and then I'll go to Steve and Astro. First and foremost, especially if we're talking about engineers, I think, and I know a number of my colleagues also believe, because we have a lot of conversations about this, that one of the most important things we want to teach our students is how to choose what to work on. Because you're going to go on and continue to learn and continue to add to your knowledge.

But being able to make decisions about, this is something that is something good for me to invest my time in, this is something that is not, these are the implications of this particular technology that I might develop, here's the ethical implications, here's societal implications, here's implications in terms of access and how will that affect the world around me, those are really important things to learn. And where to spend your time on. So we try to structure a lot of our classes to bring those elements in. That's the first thing. Second thing, I actually do think it's important for students, especially now, to be able to work in teams, interdisciplinary teams, and leverage each other's expertise, and also be really good and effective at communicating with each other and handing off tasks and being able to pass off knowledge in that way. So that's another thing that we try to bring into our classes. And the third thing that I will say, which is just sort of very broad and general, and this came up in an earlier discussion we were having today, is I think many people would agree, and I might be biased in this, but many people would agree that there really is sort of a revolution happening within the field of biology. Especially with advances that we see in genetics, genomics, the way that we treat DNA as information. And I think traditionally biology has been almost sort of a niche. Many of us-- many engineers only take biology-- don't have to take a biology class.

We might have to take a physics class, we might have to take a chemistry class, but there's not a broad even basic literacy in biology. And I think we really need to change that and treat it as something that everybody should be broadly literate in to really be able to move into this next realm of technological advances that we're going to see. And that will also allow us to think about how we bridge biology with computer science, with information science, with mechanics, et cetera, and allow us to move across these disciplines. And I think that's very important and something that is a deficiency right now in how we train students. So Astro, do you agree with her? I think that that's exactly the right set of stuff that we look for. How great are they at communicating? How great are they at learning? How practiced are they at working in groups? So that's the stuff that we should be training people to do. I'm actually not a huge believer that you have to pick what it is you're going to be an expert at now, study that really hard, and then go out and shop that expertise for the rest of your life. The bad news is that the stuff that you're learning now is going to be fairly irrelevant in 10 years. Seriously. I mean, just get prepared for that.

It is going to be irrelevant. Remember that when you're doing your homework sets tonight. But the good news is that the skill of learning things quickly, figuring out how to understand first principles and be able to reconstruct your knowledge, even after you forget 90% of it later, those skills are critical for the rest of your life. And picking up those through those homework sets matters actually way more than the actual knowledge that you're taking in. I just had someone sit down with me who used to be a tank battalion commander. His first job after being a tank battalion commander was to come to X, and he was doing field ops for Project Loon, the stratospheric balloon project. Now, a few years later, he is literally designing the balloons. He had no background in any of that before, but he's fantastic at it. And he was saying, well, maybe it's getting time to hand that over. My team is really developed and I'd like to kind of go back to the early stages and do some more stuff.

And I said, why don't you just join another field ops team? That worked for you the first time. Just go get some mud on your boots. Jump into a team. Even though you're really expert at these things now, they won't know it. And you're probably going to develop a new set of expertise that has nothing to do with balloon design by the time you're done. So don't get wound up about it. Just take that quick step back, and you'll rocket forward again. And I think the same is true for everyone in this room, that if you plan for static stability, you're going to be really frustrated. But if you can build the skills of dynamic stability, it's going to be awesome. Steve? I want to amplify that, and totally agree.

A framework for what Astro's saying is if you believe, and I think you should if you're close to science and engineering, that we are in an era of ever accelerating technological change, you can think of Moore's law as one epitome of that. Then a corollary is the lifespan of any given idea, company, technology, product, you name it, is going to become more and more ephemeral. And so your career, your life of lifelong learning becomes all the more imperative. That process learning, how can I learn, how can I engineer better, how can I pivot to something new from what I've done, are critical skills. So translating to what that means here today, think about methodologies of-- like you had the scientific method. You have methodology of traditional engineering. You have these unusual things like deep learning, machine learning, that are closer to an evolutionary algorithm. And learning more about that in biology and how it infuses all fields will be important. So if I could say what Christina said, I'd say everyone here, no matter what you're starting, to learn a little bit about biological systems and how they work, and a little bit about computer science, will be almost like the connective tissue between almost anything you want to do in your life. Because I think those are the metaphors, the frameworks that will apply to almost everything.

And then you might think, in terms of what you prioritize, what do you uniquely bring to the table? What quirky background, either academically or what have you, can you combine in ways to have insights about solutions to problems, or areas of the

world to focus on, that everyone else in the room isn't also doing? Because it's those interdisciplinary pairings of ideas, I think, where true meaningful breakthroughs come through. And that's, frankly, I think where you learn. I think that's where a lot of excitement comes from, and that's what pushes the world forward. Can I-- Please, yeah. --just give two super quick examples, that just made me think. The sort of COO for Google X, whose name is Tom, who has a lot of the central teams under him, these are people like strategic sourcing. These are people who go negotiate with our supply chain and say, hey, we'll give you a million dollars if you'll make this totally weird thing that you've actually never made before, and then maybe if we need it, we'll get thousands of them from you. That person, his background is not finance. He didn't work at McKinsey. He was first in his class in computer science at Oxford.

And the person who was running-- we have the very wide part of our funnel, which is called the rapid eval team, rapid evaluation team, which is all these sort of technical polymaths, when they graduate something from their part of the organization, when they fail to kill it for the first couple of months, the place that it goes next, which is called the foundry, is run by a woman named Obie. And Obie has no technical background. And yet all of the things that we're doing, she actually was running marketing for about a third of Google before she came to X. And now she's running a bunch of teams that are doing incredibly technical things, and she is perfect for that. You don't know where you're going to end up. And I think it's these weird sort of confluences of the skills that you're picking up and the experiences you're picking up and the style that you have that will end up determining where your best niches are in life, not what your major is. One last thing. Speaking-- they both said it. Practice it. It's an acquired skill.

Everyone can be a good speaker. It's terrifying the first 10 times. Just do it. You'll get better. The other skill I might emphasize is respect for cognitive diversity. Some geeks like myself thought there is a optimistic, high-energy, analytical way of thinking about the world, whatever my cognitive style was, and that's clearly right and everything else is deviant. And in my first few jobs, I really had a tough time with other ways of thinking. Those people were all just weirdos. And once I started to embrace that a team of five people plus or minus two that are very different is going to be better than a team of five experts that think just like myself. Thank you for that wonderful way of describing the importance of diversity.

But as I listen to you talk, I get a little worried that maybe domain expertise isn't really needed. You're really talking about those broader skills, learning how to learn. But where does domain expertise fit into this? Christina, you want to take a shot at that? I would say that I think there is value in domain expertise. And that-- but really what it is is that, from the context of engineering innovation, I believe that you have to have a good understanding of the underlying substrate that you're working with in order to apply and in order to innovate around it. I think the point that's been made, which is also, I would agree with, is that the specifics, what is sort of cutting edge now is going to be obsolete 10 years from now. And so you have to be able to continue to learn. Because if you just think about biology, our models of things 25 years ago, or how I was taught, genetics, is very different than what it would be like in a genetics class now. Just our understanding of that substrate is very different, and the way that we understand it and the way that we interface with it. And so if you want to continue to sort of be at the forefront of a field, you have to be able, you have to have the plasticity to continue to learn and to continue to adapt. And also something that I think is challenging for a lot of people that was sort of highlighted here is being able to move between disciplines, and not necessarily only putting yourself into one box but being able to move fluidly between them.

Do you both agree with her? And then I'd summarize, at least in my mind, that I think process learning and process knowledge is going to become more and more important than product knowledge. That could be instantiated in a database. What do we know about the world? What are the facts of the world? And the process learning, how do we accumulate wisdom? How do we build complex systems? How do I learn over the course of a career? How do I succeed in building a great team? That kind of iterative learning, focusing on the process variables that correlate with success, I think will be increasingly the locus of learning. Interesting. Yeah, I would say maybe, ballpark, 60% percent of the people we hire, we have some domain expertise that we need them to have. If we're building a contact lens, and we need electrochemistry work done, we're not going to hire someone who doesn't have an electrochemical background. That doesn't make any sense. But there are lots of people who have electrochemistry as part of their tool kit. That really is like table stakes even for those jobs. And there are people, some of whom don't actually have electrochemistry as a major and still have that as their expertise.

And all told, that's maybe 60% of who we hire. There's a lot of, Sally is amazing. Let's just get Sally, because we're crazy if we don't have Sally as part of our team. And if that's 40% of everybody and there's some of the other 60% where you kind of didn't learn soccer from the book, you just ended up playing a lot of soccer, whatever, you know how to play soccer at the end of the day, and that's what we need. Great, you're hired. I think people focusing on the domains for any reason other than because you love it is probably not going to get you what you want. So I'm gonna just ask one more question to the three of you and then we're going to open it up to the audience. So I just want to put the audience on notice that in a minute I'll turn it over to you. But I had one last question. A lot of the technologies, Steve, that you fund, Astro that you develop, Christina, the potential technology you develop in your lab, there are a lot of societal implications.

And some of them are not always positive. How responsible do you feel for the societal implications of what you invent and

invest in, and how do you think about linking technology, innovation, and societal responsibility? And you want to start it off, Steve? Sure. That's a deep question. So in the short and immediate term, you can be responsible by saying there are certain categories you won't invest in that are distasteful. But I think the deeper question is dangerously powerful technologies, which are all around us and compounding dramatically, as you know. And soon, the ultimate vector is heading towards technologies that will be trivial to execute as an undergrad that could kill much of humanity for almost no money. In a world-- that's a difficult world to imagine in the future. So we as humans are pretty static in terms of who we are. Our culture evolves at a less glacial pace, but still somewhat glacial compared to all these technology advances. And so there could be a bit of a disconnect, both in terms of capabilities to do harm, as well as all kinds of ripples through society.

So I believe technology is synonymous with progress. It's kind like, it's hard to detach progress in any sense from advancement of ideas. A technology could be something like the scientific method itself, or a new form of governance. That's a form of technology. But most of what we focus on are the technologies that are accelerating dramatically. And I think it's a net positive on almost all fronts. I think there are some really big implications, though, as it renders the fabric of society. So we already see how we live very differently today than 50 years ago. Culturally, versus 2000 years ago, it's astounding how much progress we've made. And violence keeps declining, and our future might be right.

But I do worry about things like an ever-accelerating rich-poor gap that might be an inevitable byproduct of network economies being part of most information businesses and most businesses becoming information-centric at their core at different rates over time, where there just won't be jobs in agriculture. There won't be jobs driving cars, which is about 20% of paid employment today. That's all clearly going away in the near term. And all of those new jobs come, and we'll all become information workers and entertainment workers. And that sounds great, except for people who don't want to do that. And it's all globally competitive more than regionally competitive. And so none of this will happen right away, but the I think inexorable trend is that we need to understand a world where that will just keep accelerating and not self-rectifying, but for perhaps policy, but for some waves in philanthropy which I find encouraging. But there's also a lot of self-reinforcing feedback loops in that. So that's one thing I worry about. I don't know, though, what as an individual and an individual technology area could do about it other than acknowledge that larger trend that's an emergent phenomenon and say, well, there's a big problem I could try to solve as an entrepreneur.

Let me think about that future, and let's say, wow, you want to provide for basic human needs for everyone when there are no jobs. Well, how about Maslow's hierarchy of needs. Food, shelter, clothing, health care and education, if Maslow was alive today, and let me do an app to provide free health care forever for everyone via a smartphone, because I think in the future we have to have that or we'll be in trouble. It can be an entrepreneurial sort of impetus to solve the problem instead of wringing our hands about it. And so that's what I'm looking for as an investor, to say technology is going to do what it does. That's exogenous to us. We're like the vessels for ideas and like a genetic sort of bag, if you will. We're also a medic bag and we spread ideas promiscuously. And that's just going to happen, and we can't stop it. So how about tilting our efforts to anticipating that future in helping the transition be a less cumbersome as it currently seems to be? Great.

Christina? So as somebody who develops technologies and thinks about technologies we want to develop, I think it's definitely the case that as engineers we have a responsibility to think about the broader societal applications of the technologies we develop. This comes back to-- part of it comes back to how do we choose what to work on. I think it's also the case, and I run up against this a lot, and I think we have to acknowledge that there's dual-use aspects of almost any technology. I mean-- and many times in ways that we can't even imagine. A great example-- I mean, did we envision or imagine that airplanes would be used to fly into buildings? I mean, and so you have to-- and so you're not going to be able to envision or even predict all the ways that a technology might evolve and all the ways that someone might take a technology. I think there are-- with many technologies, you can think about sort of short term or immediate, how it might be misused, and then you can come to some conclusion about whether you think what is-- basically, what are the pros and what are the cons, what are the trade-offs, and is the overall good worth the risk? And then I think what plays into it also is you want to think about policy around it. For some potential risk, for some potential dual use, there's a lot that we can do working with regulators, working with government, to think about policy that can be developed to ensure that as new technologies are introduced, we're mitigating what's happening in other societies and making sure that that's been minimized. And thinking again about the culture of how we train our students so that we're training cadres of students and engineers that are going to go forth and innovate and apply things positively. Great. Astro, the final word on that topic.

When we're working as scientists, as inventors, as entrepreneurs in any area, obviously there's a significant responsibility. But that responsibility can't be to control the outcome of society. That's not the job of any of those people. It's not necessarily anyone's job, but if it is anyone's job, it's the job of the public sector, not the private sector. Now, that doesn't waive all responsibility from those people. But it means that their participation in the process is one of transparency and education. That means participating with the public sector by saying, here is what we're working on. Here is what we know about what it's likely to be in the future. You can help, or we can help you think through what we might not know and maybe discover some of those things ahead of time. Let's talk about that.

I want to make sure you, the public sector, understand as best as possible, and then that's your job, not our job, to control this thing. Because when we pursue any kind of technology from a perspective of fear and secrecy, we are crippling it. We are turning it into something less than what it otherwise could have been. And we are ceding the real development of that thing to the very people we're most afraid of. So instead of painting these technologies as the poster children for our fears, we should be accepting that the world is changing, and doing our job as educators to the public sector of the ramifications of these things, but then continuing to move these things forward without embarrassment. Great. All right. Let me now turn it over, and I think your hand was first up, so you get the first question. So there's this interesting idea that generational privacy norms have developed as a result of say, Facebook, sharing photos like you were talking about before. And I wanted to ask you, to what extent do you think technology drives our value system? And to what extent are the technologies that you develop and fund a product of the values you would like to foster? And I'm going to ask you to repeat-- was that directed at Steve? All three, OK.

So to keep it going, we'll just have one panelist answer so that we can get lots of questions. Oh dear. So maybe Steve. But I'm going to ask you to repeat the question, because he's not miked and this is being taped. OK, so working backwards, the part I heard at the end is to what extent does a value system infuse the investment process we might go through, and also how does technology adjust perhaps our values over time? Was that the gist of the first part of it? Well, the second part, let me start there because it seems a bit more direct. There are absolutely, as I alluded to earlier, technologies that we think are for the good of the world. And it's not just that we feel better about our careers. It also tends to correlate with investment success that if the company you invest in is doing good, a lot of good byproducts follow, like other people like the company, employees like the company, customers like it, as opposed to a tobacco company or guns or something. So we tend not to invest in guns. There's a long list of things we tend not to invest in.

But almost anything that you'd find in Vegas, for example. That's a simple category. But the deeper thing is that I've been struck by how some of the most successful companies are purpose-driven companies. So Planet Labs was mentioned. Tesla, in an opening statement of the Q2 report last year, said profits are not a priority as the opening line of the quarterly report. And the point was there's a higher mission to help usher in an era of electric mobility. And as that as the primary goal, there's a subgoal, which is, you know, we're not going to lead the world to a new era if we're not ourselves profitable and an attractive business. So there are-- and there are many others like this. Google arguably is a company like this. Facebook arguably, connecting the entire world is the mission, and oh yeah, begrudgingly some revenue and profits might have to come along when Zuckerberg realizes it's important to carrying on.

So my point and sort of glib recasting of that is that finding a passionate entrepreneur that wants to change the world for the better, and that's a value judgment, is exactly what we should be doing. And that's also what I try to do in my career, is not just try to maximize revenue or investment potential anyone, but what is going to push the ball forward over a longer arc of time, which seems to make a more resilient strategy. I think our cultural norms and mores do shift over time. I think our circle of empathy continues to expand. And you can look at entrepreneurship itself as a bellwether of that, which is people used to think, 20 years ago, about US companies in the US serving US companies for US benefits. And now it's clearly global. In the case of Elon Musk, even interplanetary. Let's help out those future Martians, I mean, literally. That scope of the mission is amazing. And I think in the long arc of history, like the family unit, the tribe, the city-state, the nation-state, the globe, all sentient beings, that seems to be how our scope of empathy expands.

And I think it's largely by awareness, communication technologies that enable this awareness, and frankly the cultural revolution that co-evolves with our technological evolution. So again, 2000 years ago, rape, genocide and slavery were all fine according to mainstream opinion. Pretty weird. In certain cases, like warfare, when you need to smote somebody, so when smoting is on the menu, then those things are fine. So again, I'm glib about it, but 2000 years ago, humans were the same. Genetically, we're no different, except for some teeth and little things that have nothing to do with our actual brains. It's cultural evolution. And I think it's largely infused by the inventions of democracy, the scientific method, how do we accumulate ideas over time and discard ones that are bad? That's a form of technology. So maybe the scientific method more than anything else has accelerated the pace of learning progress and betterment. OK, next question.

Young lady in the front row. So you talked a lot about how the process of education is becoming more and more important than maybe domain knowledge for most people. But a lot of students when they're graduating don't really know where to find that, so maybe end up in a job like or banking or somewhere where they pitch a lot of but maybe you don't get as much of that as you seem to think. If you were graduating school now, or advising someone who were, where would you recommend going to get that process education? Astro, you want to take that one? And repeat the question, please. I will. I feel-- the question is, if we're not supposed to learn a skill, like computer science, we're supposed to learn these soft, more transportable skills that we've been mentioning up here, but that's not exactly what you're taught, and it doesn't say that on your resume the day that you graduate, where are you supposed to go both to build those skills or maybe to pick them up in the first place? That's your question. The advice I give to people, independent of age, is find somewhere that you can be passionate about. Surround yourself with people who are at least as wonderful as you are, who are smarter than you, and who you're confident you can learn from. More than anything else, find a manager who can be your mentor. I don't care if they're making smelly soaps, or

trying to go to Mars, or something in between.

If you maximize for those things, where you're around amazing human beings trying to work really hard to do something good for the rest of humanity, it's all going to work out. You can't not learn around those people. You will inevitably learn the skills that you need. And if you're in a place where the mission and the types of people you're around are draining your soul, how are you going to end up learning process points like how to work in a group? You're going to try to avoid the people you're around. So I think that the things that we wish were true happen to line up with what's in your best interest in this case, which is to really be around great people. OK, next question. Which means Google, clearly. Just for example. Yeah? So Mr. Tyler, what you said about the responsibility for the private sector in terms of the societal implications of your actions, the first thing of all three of you said was, what you need to do is you need to figure out how to change society.

That's what technology fundamentally does. So if that is your goal, to change society, aren't you a little responsible for how society changes as a result of what you do? Doesn't that put more of a pressure on you to see the impacts and deal with them as they come up? I think-- so the question is, if it's our goal to change society or the world in some way, however well-intentioned our goal to change the world is, then doesn't that add some incremental responsibility to us in how the world ends up changing as a result of what we build? I think that's a totally fair question. I'll make two points. First of all, while I certainly wouldn't use it to dodge the responsibility, if someone is making a technology, has no idea what it's going to do in the world and is just going to throw it out there without even a goal to make the world better in some particular way, I can't see how that's better from a moral perspective than actually seeking to make the world a better place. So if you pick something like, let's say, cars that can drive themselves, you need to work out early on whether you believe in your heart, collectively, that that's actually going to make the world a better place. Now, no technology, as she was pointing out, is an absolute positive. All technologies are net positives, or at least have the potential to be net positives. But none of them are absolute positives. So you have to go through the math and say, is this a reasonable thing to do? By the way, we probably throw out 30% of our ideas purely on the basis that it's a zero-sum game. That we could make money-- that's it's not even necessarily bad for the world, but that it's not strictly adding value to the world.

It's just shifting the value from someone else to us. And even that leaves us cold, and we say, nope, not going to do that one. So we go through that process. But, for example, if, just to use self-driving cars, there are going to be people who have new jobs because cars can drive themselves, surely there are going to be some people who don't have jobs anymore because cars can drive themselves. I'm compassionate about that. But A, I'm not sure that that justifies 1.2 million people a year dying in car accidents that we would just maintain the status quo of those jobs. Because the fundamental problem, I believe, is that we are failing the young people of our societies around the world. We are not doing a good job of training them. And let's say in the limit we were doing an abysmal job, and it might not be much better than that, frankly, of training the young people of the world to be really ready for the future. What should we do, stop all forward progress? Not cure any more diseases? Allow cars to be as inefficient as they still are just to make up for the public sector not being able to educate people? That seems like the wrong solution to the problem.

So I accept some responsibility. But I think the public sector and society at large has to meet technology's development halfway and take some responsibility, which it frankly does not do a great job of doing, but throws it back at the doorstep of the changes that technology causes, even when those technologies are aimed in this particular way, I think in an incredibly well-intentioned way. Maybe one last question over there. Ray Kurzweil said that maybe you have in 2001 at the latest couple years from now, and my question is, is that the end of education as you know it today? Because we can just copy-paste-- So, can we try repeating that a little bit more slowly? And it's going towards Christina. So Ray Kurzweil say that we will have digital immortality-- Digital immortality. Digital immortality by Ray Kurzweil. Yeah, so at the latest 30 years from now. Right. Maybe even 15 years. And my question is, is that the end of education as we know it, because in 15 years we can just copy and paste the skills that we get from the university .

Is it the end of education as we know it if we can upload our consciousness to a computer? And why do we need to stop learning, then? Yeah. No, wait, why would that follow? --copy/paste the skills-- But you copy paste ignorance. I mean, how do you-- How do you make that progress thing happen, of learning? Because for example-- for example, skill to bike, or to drive or something like that, it's actually-- we can see from the brain how we can do it. You don't have to-- you can do everything online. You can just download. You don't need to go to school. You just read it off the computer. So it's basically, won't online education take over everything? No. Is that what you're saying? No. No, no, just download it in your brain like "The Matrix." Oh, that kind of thing like "The Matrix." Have you got that one figured out? I know jiu-jitsu.

Like, reality check. Just because Ray Kurzweil believes in transhumanism does not believe that anybody else believes in that. Right. Just-- you're not wasting your time going to school, and I would be willing to bet your grandkids are not going to be wasting their time going to school. Can I-- I think it's actually an important reason why that intuition follows. When you grow a brain, a real brain, or you create an equivalently complex neural network or deep learning algorithm, the actual inner workings of it are completely inscrutable. It's only really understandable in its interfaces having to do with the way in which those



algorithms work. Same is true for evolution and the path dependence of how we get to where we are. So you can't-- even though Jeff Hawkins at the end of his book on intelligence just out of nowhere says we'll cut and paste, like, French language skills between these evolved artifacts, I don't think that's possible at all. Because we have no example around us in information networks, rich complex information networks, where the evolved product of biological evolution, deep learning, machine learning, any of these, is cut and pastable.

What is much more efficient is just grow a new one from scratch. And so you'll be learning more than ever before, just more rapidly. OK. We are at the end of our time. I want to thank people for coming. And I especially want to thank our speakers for their very interesting insights.