



Transcript: AI, Automation, & the Workforce—Surveying the Future of Work

On Tuesday, April 23rd, 2024 the Center for AI Policy held a briefing for House and Senate staff on AI, Automation, & the Workforce—Surveying the Future of Work.

The speakers Robin Hanson (George Mason University) and Simon Johnson (MIT), moderated by Jason Green-Lowe (Center for AI Policy), discussed AI's effects on jobs, ways to help workers adapt, lessons from previous disruptions, and more.

What follows is a full transcript of the discussion. The transcript was generated safely by AI with human oversight. It may contain errors.

To watch a corresponding video recording, see [this video](#) on the [Center for AI Policy's YouTube channel](#).

Full Transcript

Jason Green-Lowe | 00:07.552

All right. Thank you, everyone, for joining us. Welcome to the Center for AI Policy's briefing on AI in the Workforce. I'm Jason Greenlow, our executive director. We're a nonprofit, nonpartisan organization that's working with Congress to help make what we see as the enormous transition to the age of artificial intelligence as smooth and as safe and as productive as possible. I want to open with this short advertisement for Cadbury Chocolate. I promise we're not affiliated with them in any way. We're not taking making their money. I just think it touches on some relevant themes.

Cadbury Ad | 00:41.756

Looking to upskill for the future? This new AI can perform all coding jobs in seconds, including blockchain development. While this AI is already outperforming accounting firms. Meanwhile, the new graphic design AI aims to automate graphic design and could minimize the need for...

Relax.

Because in a future where AI does most of the work, there'll be one thing that humans will finally get to do all day long: nothing.

So it's time to learn the most essential skill of the future. Five Star presents Nothing University with the first ever diploma program in doing nothing. An intensive course covering the most future ready skills out there. Well, let me check with Ian and get back to you in a few days. I mean, With our state-of-the-art workshops, campus placements, and guest lectures, the world can finally prepare for the glorious future that awaits us. Eat five-star, do nothing. Enroll now.

Jason Green-Lowe | 02:05.299

All right. So the ad is funny, especially if you enjoy a certain kind of corny melodrama, but it touches on a serious anxiety, which is what kind of work will there be in a future shaped by AI? So to help us answer that question, the Center for AI Policy has brought in two distinguished economists. Robin Hanson, immediately on my left, is a professor of economics at George Mason University. He has a PhD in social science from Caltech, a master in physics and philosophy from UChicago, and worked for nine years in artificial intelligence as a research programmer at Lockheed and at NASA. He's the principal architect of DARPA's policy analysis market, and he's developed new technologies for conditional, combinatorial, and Intermediated trading. In 2016, Oxford University Press published his book, *The Age of M, Work, Love, and Life When Robots Rule the Earth*, which studies the societal consequences that could follow the development of the first truly intelligent digital minds. A little further on my left, Simon Johnson is a professor of entrepreneurship at the MIT Sloan School of Management, where he is head of the Global Economics and Management Group. In 2007 and 2008, he was the chief economist at the International Monetary Fund. He currently co-chairs the CFA Institute's Systemic Risk Council. In February 2021, Professor Johnson joined the board of directors of Fannie Mae. His most recent book, *Power and Progress, Our Thousand-Year Struggle Over Technology and Prosperity*, explores the history and economics of major technological transformations up to and including the latest developments in AI. I was going to bring a copy of Robin Hanson's book, but turns out I read it digitally, which perhaps is appropriate given the themes. I'm going to lead these two professors in a directed discussion on AI in the workforce. So what is that? We're trying to provoke an interesting conversation that's not a debate because they're not opponents and they're not trying to prove each other wrong. But unlike a standard panel, we're going to let them talk to each other directly, respond to each other's points without me interrupting them every single time. I might interrupt them once in a while to steer the conversation toward a more fruitful or interesting subtopic. And we hope that their conversation will highlight. get some contrasting views on what effects AI will have on the workforce, how soon these effects are likely to peak, and what, if anything, Congress can or should do about that. So I have a few topics that I'll use to kick off that conversation, and then we'll solicit more topics from all of you. The first topic I want to touch on is the changing nature of work as technology continues to diffuse across the

workforce. So if we rewind, you know, a couple of decades, how How did computers and the internet change the nature of work for the average American as those technologies spread through the workplace? And we can move forward through history from there.

Simon Johnson | 04:57.732

Who's going first? Me?

Jason Green-Lowe | 04:58.864

I guess so, since you asked.

Simon Johnson | 05:00.902

So digitization, so roughly from 1980, did have some profound effects, certainly for some parts of the labor market. I think in terms of its overall effect on productivity, there was a bit of a jump, you remember, or a boost in the 1990s. Overall, perhaps a little disappointing compared to what expectations were, but there was something. The main impact, though... that economic research is focused on is what happened to the so-called middle skill, middle income, middle education jobs. And the job market polarization that has been much talked about and was definitely a phenomenon through the end of the 2010s as a question of whether it's receded now. But a big chunk of that, maybe 40%, maybe 80%, depending on who you look at or who you talk to, came from automation and came from technologies related to digitalization, including in manufacturing. Sorting out the precise effect is confirmed by the fact that there's a lot of globalization that goes on at the same time, the so-called China shock, which impacts very similar people in similar jobs in similar locations. So we have globalization and digitization, and those things are not, of course, unrelated. Jason, because a lot of globalization is made possible by the big fall in the cost of communication and the fall in the cost of computing, so therefore you can run a business globally in a way you couldn't in the 1960s and 1970s. So digitalization with globalization, profound impacts on the middle part of the labor market, and also boosting income opportunities at the higher end for people with a lot of education. Not so good if you're at the lower end of the labor market, in part because people who are in the middle got pushed down to compete with people at the lower end.

Robin Hanson | 06:32.468

I guess I'd like to call your attention to the longer timescale, the bigger picture of several centuries and the whole world. And the bigger picture is that for several centuries we've been in the Industrial Age, after the Industrial Revolution. And the world economy has been growing roughly steadily for a long time. Different parts grow at different speeds, and sometimes one part is speeding up compared to others, and sometimes they're

slowing down. And at different points, different parts of the economy have been doing well or better. Sometimes certain industries grow faster than others. Sometimes certain industries draw employment. Sometimes they lose employment. And over the decades, we've had new kinds of technologies that showed up. Once upon a time, just... tractors, engines, light, refrigeration, etc. If you look over the long term, we've just seen a lot of new technologies arise, and they have different effects. Sometimes the people at the top end of the wage distribution are helped the most, sometimes those at the bottom and the middle, sometimes it's urban, sometimes it's rural. The highest level story here is just that over centuries, we've had these unpredictable ways in which the economy changes, and we have a general way the economy adjusts to these things. It's called the labor market, called capitalism. Basically, people try to anticipate which way the current, you know, who are going to be the new winners next. You know, the capitalists try to invest in those firms. The workers try to learn on those jobs. People try to move to whatever cities or towns they expect to be the winners. And sometimes they're guessed right. Sometimes they're guessed wrong. And the economy has what we call churn. Basically. You can't just take a job when you're 18 and expect to have that job for the rest of your life. You have to be ready to find out that maybe your job was more of a loser than you had bet on, and you need to switch to something else. Or maybe you have to realize you have to change how you do your job, even if you get to keep that sort of job. Maybe you have to realize that maybe you should move to some other place, and that your town isn't doing so well. The fundamental thing to keep our eye on here is that we have a way our economy adapts to all these unpredictable changes. And as long as the changes we see are within the range of changes we've seen for the last century or two, we should expect to be able to handle that reasonably well because we have handled it reasonably well so far. So in my mind, the biggest question about AI is, should you expect this new burst of technology to be in the range of things we've seen? such that we will, you know, have to make guesses, have to adapt. Some people will win a bit, some people lose a bit, people will have to move around, but if it's within the usual range then let the usual mechanisms handle it. Maybe we can do a little foresight and maybe give people some warning, say, hey watch out for this, or hey look for this opportunity. Those are the kind of things we economists or academics could do, give you a little hint, or maybe just say, you know, don't count on things too much, let's learn this history about how Things can go how you can need to move, but that's the highest level point I'd say, is just should you be scared that this is going to be a big deviation from this whole history we've seen for a few centuries, or should you expect this to be in that usual reign and therefore our usual mechanisms will probably work fine?

Jason Green-Lowe | 09:46.845

I think it's a great question. What do you think about the current generation of generative AI systems? So if we're talking about ChatGPT, DALI, or maybe some slight

improvements on those systems, does that fit within the amount of churn that you see is normal or historically average? And maybe also... does it matter whether we're on the low end or the high end of what's historically average?

Simon Johnson | 10:12.669

So, I think that's a very good question, and I think the way Robin frames it is exactly right. Are we within the sort of normal parameters or normal adjustment? I think there's a sense in which, well, there's a scenario in which AI could actually be better than some of our previous disruptions, because if what is developed by the market or by other forces, government certainly has shaped technology, DARPA helped promote self-driving cars, for example. So, some of the forces that shape technology give us. Either AI co-pilots for lower skilled, less educated people, that would actually go the opposite way to what I said about digitalization. That could actually be rising productivity and rising opportunity for less educated people. And I think the really big change also, thinking about Robin's long sweep is Can people, what are the barriers to more education? So if economists say, they often do say, more education, more flexible human capital is good. At the end of the 19th century, early 20th century, when people came out of agriculture, there was a big high school movement and a lot of improved Low-cost access to more education, which was very timely and not coincidental. If AI lowers the barriers to education, and if it helps you be more flexible so you can learn things at 20, perhaps your ability to learn things at 16 will be the same, but your ability to learn things at 26 or 36 or 46 might be better, then we might actually be easier to deal with this disruption than in some of the previous disruptions where we've relied on more hierarchical, fixed, conventional things like K-12 education.

Robin Hanson | 11:41.778

I agree. That is, if you heard stories about a century or so ago, and people had to leave their small towns and go off to the city and, you know, learn whole new careers and whole new kinds of systems, it was painful to a substantial degree. It was a lot of effort that people put in to change their lives, to adapt to changing technologies and conditions. And... We're richer now. We have a lot more tools to help us with this. So these adjustments should be easier now. We have computers to help us retrain. We've got, you know, ways to move more easily from one place to another. It actually should be a bit easier.

Simon Johnson | 12:17.894

I agree with the should. I think the caveat or the concern would be the last, back to the digitization and the globalization, because, again, presumably we face some combination of this. So it should be easier. We definitely, being richer is important. And also. having more time to think about these things and having this kind of policy

discussion or pre-policy discussion. I think the concern is that nobody set out with PCs to disrupt the labor market or destroy middle class opportunities. I really do not think anybody had that intention at all. And yet what's happened over the past 40 years has, at least through the end of the 2010s, was polarizing the job market. and there's a lot of anger and angst and frustration that comes out of that. Some of it rightly oriented and some of it perhaps misunderstanding what had happened. But still, I think that job market polarization, you know, we should have done better in the 80s, 90s, and 2000s. We should have done better than we did, given what we knew then, given how rich we were then. So hopefully, Robin's right, that we can take advantage of these opportunities, and hopefully we can have the market do as much as possible. But there is this concern given our starting point.

Robin Hanson | 13:28.394

But I think the thing that most animates the discussion at AI is this idea we could have a lot more disruption than we've seen in the past. So certainly over decades, we continue to see news articles about how huge fractions of the economy might suddenly become unemployed due to automation. And this is just a minute repeatedly. There was actually a presidential commission in the 1960s about whether the new computers would automate most jobs. I got into AI in 1984 when there were lots of news articles about AI might be automating most jobs then. You know, so basically this is I think the core emotion of automation. It's not that it might be a modest disruption or maybe we wouldn't quite track it right. It's more that could this time be really different where huge amounts of automation happen in a short time and then we would need to do something different than we've done in the past. And for that I would basically say well so far this disruption hasn't been very much at all. That is generative AI is exciting but it hasn't actually changed that many jobs. Maybe it will. I would say my fundamental recommendation is just to track. these things. That is, as you see more jobs changed and displaced, we have stats that we used to collect. I've done stat analysis on jobs in the US and their degree of automation in each year and what that correlates with. And my analysis from 1999 to 2019 was that there was basically no change in that time in the rate of automation or the types of automation. But we should be tracking that. And to the extent that we could, you know, see that a trend was deviating, that is something was different from the past, we could start to call attention to that and do something about it. But it hasn't happened yet. And we've got all of these people crying wolf for decades saying it's about to happen when it doesn't. So just wait and see if the data warns you of something. But Stop listening to all these just warnings on the basis of really pretty little.

Simon Johnson | 15:21.650

I agree. And I think there's also rhetoric from the business sector that is exaggerated and sounds good to invest and so on. But I would like to highlight also, in addition to concerns about automation, potential concerns, which may be exaggerated, I think the potential for augmentation is really interesting. So are we developing, are the people driving the technologies, are the people who are hiring the... the best engineers, are they dry? Is it their priority or particularly interesting to them to try and boost the productivity of those relatively low productivity workers now? So develop the co-pilots for electricians or plumbers or other modern craft workers, nurse practitioners, for example. And I think to the extent that they can be, to the extent that we think that's a good idea, to the extent we think augmentation would be helpful, and to the extent we think the market might not provide it because you don't make that much money. or changing education is particularly, you know, politically difficult, or whatever reason, then I think sort of initiatives like the DARPA self-driving car, but applied to augmenting technology, doesn't cost a lot of money, gets a lot of attention, sort of motivates engineers, that is, that's extremely interesting to think about.

Jason Green-Lowe | 16:29.488

I want to dig into this issue of the retraining cycle or the educational cycle. You both raise a great point that, okay, we have more resources than we used to that could, in theory, be put into education. Are we going to? How do we know if we're going to? How many years does it take to retrain a worker on average? And how many years does it take for a job skill to become obsolete? And how do we think about which cycle will happen faster?

Robin Hanson | 16:56.722

Well... I'm gonna have to, you know, mention a sad truth here, which is most of, like, college education isn't actually very useful on jobs. So but we do it to show that we're better than other people. So a lot of education is just showing off. showing off how smart, conscientious, conformist you are. It's a function of our society to help us sort who's better than who, but it doesn't actually teach us much about any job. So that means there is a huge opportunity then that we could divert some of the resources we spend on not learning into learning other things. That would be good just regardless of AI, but it might be especially good in the light of people needing to retrain.

Simon Johnson | 17:41.732

So I think Jason, one good question always is which Neil Stevenson novel do we actually live in? And I think the Diamond Age, which he wrote at the end of the 1990s, is an interesting one to look at, because in that, without giving too much away, in that education becomes much more decentralized through actually devices. This was he's writing the late 1990s. These are devices which empower people, particularly young

women in China, is the focus in that novel, to learn and to be able to figure out what they want to learn. And I think that. Reorientation about thinking about what is education, what are skills, what am I missing? One of my colleagues in the mechanical engineering department at MIT, makes this recommendation to all of us, which is get the best latest AI, not the free one, the one that you pay for, GPT for, for example, and use it to teach yourself something that you don't know and that has a bit of a hurdle, like a programming language, for example, and then apply that new knowledge to a problem you know very well and something you do work on all the time. And I think the problem that a lot of us have, particularly busy professionals like people in this room, is just don't have one, two, three days. or a week to do this new thing, get over that hurdle and then see to what extent you can apply it. And I think that's just one microcosm example of what does it mean to expand skills in a way that's useful to you in the tasks that you already do, you already have assigned to you.

Jason Green-Lowe | 19:07.379

I want to press both of you again to see if we can put some numbers on all this. I take your point that the new technology could allow you to learn faster. I like Duolingo. I like getting coached by GPT or whatever. And I take your point that much of college is about signaling. But how long does it take to learn a new job skill if you go and focus on it? How does that compare to how long it took in the past? How does that compare? A lot of your book, *The Age of M*, is about... decreases in the doubling time, right? This idea that we're sort of doubling the economy faster as time goes by, less time between new major innovations that kind of change the paradigm of how people work. So is education also speeding up in practice in terms of how quickly people can learn new skills?

Robin Hanson | 19:57.554

Here's the highest level principle to keep in mind about. change and regulation, which is the less you know which way things are changing, the more you want to be flexible and allow people to adapt to whatever comes. And so the more you try to prejudge and pre-structure and regulate what people are allowed to do or how they're supposed to do it, the less flexible they can be. So I would say... let's not like make new programs where we tell people how they're supposed to learn this stuff. Let's encourage people and empower them to try lots of things out and let's find out what works.

Simon Johnson | 20:31.668

Right. Yeah. So developing tools, I think, Jason, I mean, is education speeding up? I think the answer is not yet. And no, and it's frustrating because we see many things changing around us in the world. And including, of course, with AI, we see jobs that were or tasks that were previously assigned to pretty well qualified white collar workers. Okay, there's some status in the in the in the well qualified piece there, but that's You do

need some other kinds of skills that these people have acquired. These tasks can now be done by a machine, by an algorithm, and that's you know relatively Gets your attention, let's say. But I think the right answer, and I think this is in line with what Robin's saying, but he can speak for himself, is that if you give people the tools and show them how to learn new things and convince them that learning doesn't stop when you're 18 or 21, and actually it's relatively easy, you can gamify it, for example, you can have ways that multiple generations work together. So the younger people have some skills, the older people have some skills, and you're moving on this in parallel. I think we can speed up useful education. The development of skills that actually are productive, improve our lives, and deliver solutions that we want.

Robin Hanson | 21:45.546

I think we're over-credentialed. That is, we rely a bit too much on credentials in our society, and that's a bit of an obstacle here, because whoever makes the credentials isn't going to adapt that fast to changing what an electrical engineering degree is, etc., on the basis of the new technologies available. So that means people will often be learning things that are useful but aren't credentialed. And that will be an obstacle here because people will find it hard to use their skills that they can actually acquire through new means if the existing credential systems don't give them a path to that.

Simon Johnson | 22:21.433

One thing that concerns us is at MIT, almost half, I think it's 48% of undergraduates are now majoring in software, in something related to AI. And I think there will be much better. And so there's a crowding effect, there's a hype effect. Much better would be to go into other things that are neglected, including mechanical engineering and other forms of... hands-on engineering hardware, for example, but have AI as a tool, as a co-pilot, boosting your ability to solve problems there. I think that's a really undervalued area. But to Robin's point, you can't tell 18-year-olds what to major in. But you can't say, we're not going to let... You can only come to MIT if you promise to do mechanical engineering and we'll expel you if you don't. That's not going to work, right? So I think having these conversations, showing people the opportunities, pushing in various... in various directions. I testified recently to the Senate Armed Services Committee about the defense industrial base and workforce issues, and there was a lot of really interesting bipartisan commentary from the senators about the need for these practical hands-on skills, the need for welders to become more productive, because that's a key constraint in the building of submarines, for example. I think that's a very... And then, so why don't we have enough people going to welding? Is it the pay? Which I think it probably is in part. Is it other things about the conditions? Is it how people perceive opportunity? What can you do to communicate? There's a career, there's lifetime opportunities, there's affordable housing close to where we build submarines, for

example. Whatever the constraint is, I think addressing that head on is the right thing to do.

Jason Green-Lowe | 23:48.821

Is there a point at which adding tools, adding co-pilots to make somebody's job more efficient or make it easier runs into diminishing returns? I mean, maybe there's more to be done on welding, but as an office worker, I've already got spellcheck, autocorrect, pre-formatted templates. I've got spreadsheets. I've got online search that gives me a whole wall full of encyclopedias at my fingertips. My email will suggest a reply for me. How much easier can my job get?

Robin Hanson | 24:16.480

A lot.

Jason Green-Lowe | 24:17.585

Yeah, tell me about it.

Robin Hanson | 24:19.114

But just the history of the economy over centuries is that that's how it's always looked to people doing what they do. It's hard to see all the ways that things can be done better. It's hard to envision them. And, for example, I have some colleagues who basically say, look at the changes happened in the entire last century. There's just no way we could have as much change in the next century as that. They just literally can't imagine it. They see so many fundamental things in our lives that have changed in a century and they can itemize those and they can say, look, those have changed about as much as they could possibly change. So therefore they predict that the next century will just have a lot less change because they literally can't imagine it. But that's how it's looked every century for a long time. It's just hard to imagine what becomes possible. until you do it. So that's why people should just be experimenting, searching, and trying things. And to your point about all the people going into computer science, a basic economic fact is most of the value for any new general technology is in applying it to specific things. That's where most of the value will be achieved. That's where most of the income will be achieved. Most of the money of capitalists making money will be achieved in applying the abstract things to specific things. Like... Computers remade the economy over the last 30 years, but the computer industry is still really small compared to the rest of the economy. But we have data basically that in a large number of industries, the leader in the industry was the one who found a way to use computers better than their competitors, and they got an advantage in being good at IT, and that gave them a heads up over other firms in their industry. So there was huge value achieved. by specific firms figuring out how to use computers to do that industry's job well. that's where most of the

value comes, not from the people who are the computer industry per se, and that's how it'll play out again for AI. You shouldn't be trying to get a generic AI skill per se. You should be asking, how can this be applied to something? And can I be the guy who makes this abstract thing work in this specific context?

Simon Johnson | 26:22.477

So I'll go to a little long history now, but then answer your question, Jason. So I think Alexander Field and Bob Gordon have made a very convincing case that there was a lot of innovation in the early 20th century, for example, in the US, which spread around the world. And this was tied to coming out of agriculture into cities, a lot of labor-saving devices so that women's lives were fundamentally changed and their ability to have more options, including if they want to work in the labor force, became possible. And that big, decades-long transformation drove a lot of the shared prosperity that we got after World War II. So we can talk about other conditions that were in effect in the 40s, 50s, 60s, and 70s, but there was a big, sort of a long push on technology. I think what we're looking at now, to the point Rob made earlier, is... Continued innovation definitely changes hard to predict. Not as big a change in transformation as we saw in the early 20th century, but it's still substantial. And I think that the, how much easier can your job get? Well, we don't have full data, so we do need to see more examples and we need to see actual implementation of this, but we can already see transcription for call centers, for example. A medical call center, a health insurance call center that I'm familiar with, used to take seven minutes per call. Now it takes six because the AI does the transcription. It does a better job than the human and is completely HIPAA compliant, right? So there's an immediate bump in somebody's productivity or perhaps you need fewer workers. Call centers being used more generally, um, or people in call centers are using co-pilots and reporting in at least one or two studies that it's for the less experienced, less well-prepared for the job, they're doing better, and their productivity has increased, they're more able to answer customer inquiries, and the customers are happier. So there's some effect there. There is a concern that the sort of tier one call center jobs may be more than maybe replaceable with AI. That remains to be seen, but perhaps that generates more opportunities for Tier 2 call centers, people who actually solve the harder problems. Maybe you get to them earlier, for example. I think the most interesting example, which is still a proposal, I haven't seen it actually implemented, is for a global apparel company that employs product planners who have MBAs, and they get paid about \$250,000 a year to plan exactly where production will be made and to manage that. Cycle. My colleagues working on this think they can replace 10 product planners with one super planner who will be paid a little bit more money, maybe 280,000 as opposed to 250,000. Now, if we start to see that sort of thing, so that's an enhanced productivity of one MBA, I'm not too worried about what the other MBAs are going to do. I think they will find other opportunities. But that's an interesting...

augmentation of a highly educated person. Why can't we focus more and attempt to have more augmentation of people who have less by way of credentials or formal training? So

Robin Hanson | 29:29.578

Simon will be very well of this fact, but maybe many of you aren't. The world economy doubles roughly every 15 to 20 years. And the only way it can do that is if typical industries are getting more and twice as capable over that same period. So that's the status quo. The status quo for the last century has been the world economy doubling 15 to 20 years. And that means most industries are getting that much more effective over that time period. That's the status quo. So when you see these remarkable improvements coming along, that's the thing that can make you hope that we could keep up this pace. Because often you look at... You know, improvements, you go, we've run out of improvements. What else could we do? It looks like we're the best we can do. But we've consistently found ways to double the economy over those timescales. And these are the hints that tell you how we might do it in the next 15 to 20 years.

Simon Johnson | 30:23.162

And of course, some of the angst in the US, including I think on Capitol Hill at the moment, is about other countries growing fast and other countries learning what we've learned. AI is an interesting space because it does seem that we retain some advantage. Even there's been a lot of discussion about China and China's access to data and China's ability to train AI. I don't think we've yet seen them acquire the same capabilities that the US tech companies have. But there is a race on. There's no question about that. There's no question about whether... I mean, I think on net, AI will definitely has the potential to help economies like ours, which have a lot of educated people and a lot of flexibility and pretty good market structure and I think generally speaking sensible government policy. Other countries around the world. China is a separate discussion, but other countries I think may well struggle, particularly when and if AI becomes more capable in a manufacturing capacity. It's not yet very good in factories. You need a much more digital environment to really take advantage of it. But there is the potential for much more automated manufacture of goods. And if that comes to fruition. That's a problem for countries that rely on exporting based on cheap labor, because cheap labor will no longer be the kind of advantage that it has been in recent decades.

Robin Hanson | 31:36.885

So, just worth making the point that the US has an advantage on this, as you may notice. We got lucky, perhaps, or this is a way we paid off, you know, the final reward for our years of pushing the envelope, but we're winning on this. And so, we should go as

far as we can, as fast as we can with this, so we can take advantage of that win. There is a risk that we will slow it down out of anxiety or just generic concern. then we won't have as much of an advantage.

Simon Johnson | 32:07.351

I actually don't think there's much risk of a slowdown, in part because as soon as I said China, everybody took a note. I noticed that. So I think the risk of slowdown is not very substantial. I think the risk is, and I think we will do a lot of automation because you naturally do automation. And when you come to a company or to an executive and say, right, here's this new technology, a seven minute call now takes six minute call. You're like, okay, let's figure out where else we can do that. So you're automating the process. I think what we might do a lot of or not is the augmentation. I think that is less of a salient issue in the corporate mind. I think it's actually changing a bit. Google, for example, put out a paper in January written by their tech people and the top executives saying, hey, we should think about augmentation, which is very interesting. Microsoft has a partnership with AFL-CIO. I didn't tell them to do that. They did it by themselves. Okay, that's the market, Robin. But the point of that is to try and get union voices at the table, as they call it, for thinking about, okay, what does technology, what should, what do we want to achieve with technology? So I think Robin's right, there's this big edge that we have, and when you have an edge, when you have the lead, you know, you obviously can set standards, you obviously can have a big influence on the global conversation, but you can also, we have companies that have the capability to say, right, we're going to do this, this is important, let's go in that direction. And, you know, we have plenty of government conversations with the private sector. We see it around the space program, for example, where NASA works now with private entities in a way that was not imaginable in the past. So, looking for government partnerships with the private sector where we're saying, look, here's a direction we want to go in. The private sector can't go there itself because too much capital, too much risk, or too much concern about what future policy would be. So we're making some commitments, a pre-purchase, or encouraging certain kinds of innovation that we think would otherwise be neglected. I think that's a really interesting additional possibility that we have and something we could take advantage of sooner rather than later.

Robin Hanson | 34:00.218

Jobs are a set of tasks. So, I mean, Simon, again, understands this, but I think it's important to see the world of work as really thousands of little tasks. And part of the job, part of our challenge in designing the economy, designing work, is to figure out what are the tasks that need to be done, what tasks could we do that would be helpful, and then how to package those tasks into jobs. And then when you have technology, what you do is you swap out, automate a task, which for the job looks like augmentation. That is,

augmentation on a job is where a particular task a person does is now done by the machine. So at the micro task level, you have automation. But at the job level, you could either have more automation as, you know, maybe most of the tasks of the job are displaced or. augmentation when you are automating some of the tasks. So what we want to take advantage of new technology is to be able to explore as much as possible new ways to reorganize the tasks, to decide which are done by machines and which are done by people, and to maybe invent new kinds of tasks that would be helpful. So the more flexible our system of allocating tasks to jobs and allocating jobs can be, the more we can reorganize in order to take advantage of this. So there's a famous story for electricity. ...beginning of the 1900s, which was... The introduction of electricity didn't actually give that much economic boost until factories could be reorganized to take advantage of the new opportunities electricity gave for reorganizing where motors were basically, what motors were doing. And so, and that was also true of the early days of the PC. We didn't see a lot of productivity gains until people could reorganize the jobs to take advantage of the PCs. And that's probably going to be true for AI. The boost will come not from just taking AIs and directly. substituting them on the way we're doing things now, but reorganizing the whole set of things to take advantage of these new possibilities.

Simon Johnson | 35:58.591

Yeah, so Henry Ford, when he came to Detroit, the U.S. car industry made 10,000, 20,000 cars a year, all artisanal, handcrafted. Ford put car assembly onto a moving production line, and he brought electricity in a brand new plant to that assembly line. And the productivity was massively increased. So in 1929, there were 400,000 people in the U.S. making cars or in the car industry. And Ford and GM combined made 3 million cars. And a lot of those jobs were good. Good paying jobs. And then the UAW also got involved, so we can discuss that. But I think unions were an important part of the shared prosperity in the car industry for a long time. The key thing though about those 400,000 is most of those people had tasks, were doing tasks that no human had ever performed before. So you're creating new tasks. That's very important. Many of those new tasks required expertise. So you learn something about this particular task, about the organization, about how to work with the colleagues on the team. You couldn't, you can't, if you could just hire somebody off the street and with 20 minutes training they can do the job, you never, you don't get a premium wage for that. So what you're looking for is an economy that can generate a lot of new tasks with expertise and preferably, of course, this is again what the Congress did very well is there were plenty of opportunities for people who had only finished high school or maybe hadn't finished high school. That's what you're hoping for in the... well, that's what we got in the 20th century. And that was a key part of why so many people were able to reach the middle class and have relatively secure middle class lives by the 1970s. But that's what's also

eroded or did not happen as well and as, you know, sort of the automatic version. Of the way we ran the economy from the 1980s did not deliver that so I think that's why no Robins right every technology every Form automation is augmentation for someone another what I'm adding literature references Kurt Vonnegut's play a piano His first ever novel published at the end of 1940s is a world in which there are very few engineers that run highly productive firms And everyone else has make work schemes. They're not unemployed. They are paid because the economy is so productive You can afford to pay them, but they're not very happy with this sort of two-tiered structure, so I think that The more we can do to think about and to generate new tasks, new tasks that require expertise, and allow the market or encourage the market to find those things, and then find ways to augment the less educated, people without a high initial level of skill, and allow people to acquire the skill and expertise that provides them with a premium wage, those are the pathways we need to find.

Jason Green-Lowe | 38:41.026

Well, so what do those jobs look like in the near future? Instead of an auto mechanic or an auto dealership manager, how do you become an expert with significant barriers to entry in a prompt engineering where the thing you're prompting changes every year?

Robin Hanson | 39:02.599

If we empower people to create barriers to entry based on credentials or, you know, their established position, even union position or regulation, we will limit the ability to adopt this technology rapidly and effectively. So that's, in my mind, the fundamental tradeoff here is when there's a lot of uncertainty and there needs to be a lot of change, you can't set that many constraints on how it's going to play out. Let people search, let people try things. When that starts to congeal. then if you don't like something you're seeing, then that's the point to step in and maybe, you know, limit or direct it. But until you can even see the rough outlines of where this is going, it's too early to be making those choices.

Simon Johnson | 39:41.522

Yeah, it's way too early to know exactly where the jobs will be. I don't think that the... The priority should be more prompt engineers. I think that's an elusive one to chase. I think if you believe, for example, the country is on and will remain on a clean energy path with a lot more electricity, you need a lot more electricians. So finding ways to train, not increase barriers to entry. That's not a good idea for job. But if you want a lot of...

Jason Green-Lowe | 40:07.512

Maybe it's the wrong phrase. I meant expertise.

Simon Johnson | 40:09.614

Expertise. Yeah, expertise, which is different. I think enabling more people to become electricians in a way that is... That they can be productive, they can be safe, that you can scale up whatever it is you need to do in terms of building out parts of the electricity grid or strengthening or making it more robust, how we want to characterize that. I think that it's those craft type jobs and bringing more technology to those people in a way that helps them. That's a really interesting conversation. I think you should be trying, you should, I've had this conversation with people around the world, like I just mentioned Brazil because it was recent, but I don't want to pick on the Brazilians. I don't think they're unique in this. I also had the same similar conversation in Indonesia not too long ago, which is they said, okay, how are we going to get things from AI? What is Silicon Valley going to deliver to us that'll help us reorganize Indonesian agriculture or anything about Brazil? And I think the answer is you get nothing if you just sit around waiting because Silicon Valley is not that interested in Indonesia or you know The Brazil beyond some superficial level But if you think about how to use the technology to solve the problems you want to solve Your problem is not in the problems. They were discussing in Indonesia. It would not the problems of you know AI per se, they were how do we raise the productivity of this particular category of engineers and pharma for example. How do we have more good jobs in manufacturing of any kind for example. So I think figuring out how to use technology either off the shelf or Using what's coming out of the leading edge and finding ways to adapt it to your local conditions, which are fundamentally different from the United States. But then if you apply the same point to the United States, what is the problem that you want to solve? And if you don't take the shot on goal, you don't generally score, right? So identify the problems. Problems in which technology can be brought to bear and then think about ways in which AI may or may not enhance that and think about ways in which you can stimulate the private sector to do it or make some complementary government action that encourages the private sector to move in that direction.

Jason Green-Lowe | 42:10.534

Well, I can't ask for a better straight line than that. I want to throw it open to questions from the audience, particularly if you have a problem that you want to solve that maybe could use some economic insight.

Audience Member | 42:24.541

Thank you both very much. Really interesting to learn from you today. In that 2019 book, Ghostwork, Gray and Suri coined this term, the paradox of automation's last mile. And so they're looking at the people that are doing data labeling work, interacting with these models to train them. There's a huge demand for data. And so we've seen various news stories about. People in the Philippines doing large amounts of data lately and

things like that. The basic narrative here is that these models kind of they break in various ways. You need continuing human intervention and you can procure large workforce, a large workforce of piece workers to do this work. But often these models end up, or likely will end up, replacing people who are on more permanent jobs with benefits. And so there's a dystopian vision here where you have large quantities of people doing this piece work whilst more permanent jobs get displaced. I'm wondering what, I'd be interested to hear how both of you react to that narrative of an increasing amount of humans doing data work to keep AI online as it displaces more permanent jobs.

Robin Hanson | 43:31.307

The basic fact is we just don't know which way this is going. And in the past, things have just gone unpredictable directions for different technologies in different places and times. And we'll just have to see where it actually goes. I get that you want to be afraid and say, ooh, maybe this thing could be bad. But we won't know until we try.

Simon Johnson | 43:48.469

I think there's going to be a lot of automation and I think what you're talking about is a corollary of that. And I think what you need to do is run harder on the new task creation and on the augmentation side. Robin's right, you don't know which way the balance is going to come out and maybe it'll be totally fine, but I would lean towards thinking about constructive pro-market ways to encourage augmentation. On the wage point though, on the worker point, I think there is a a looming global issue, which is many of those jobs are really, really bad jobs. So content mediation, for example, for not going to name any particular platforms, but reportedly it is absolutely horrible. And really, I don't think any of us in this room would take that job for any money, actually. So there's an interesting question about, you know, these people in low income countries want to work. This is their choice to do this for the money. But. Are we entirely comfortable with that global division of labor? Is there going to be backlash against that? That's an open question. Yeah.

Audience Member | 45:00.000

How much is the domestic sort of the hardware side for AI? Are we meeting the needs to sort of attain all these different things that we might want in an application, but actually having data centers? technological advancement on the hardware side of the people with all the innovations that we want in our pocket.

Robin Hanson | 45:22.638

We have a robust competitive hardware industry that supplies whatever the hardware demand it is. I think the key thing you're seeing lately is just there's this fear of missing

out. A lot of people jumping on the bandwagon and they're really eager to be there as fast as everyone else. And so they're bidding up the price of hardware so that they aren't second. Now, often that's a bad play because often being a good second is better than being sloppy first and it's not clear that they will actually get the returns from being that eager to be that fast as soon as everybody else is, but we'll find out. But that's a temporary thing in the sense that one way or another, this initial hype boost will pass and then we'll be more of an ordinary hardware market where if... you know, the hardware will be supplied that people demand, and they'll only be demanding it if they're getting some value from it, and then it'll be fine.

Simon Johnson | 46:12.918

Yeah, I think the Chips and Science Act was on the whole very sensible step in the right direction. I think there are a few additional pieces that you could have in terms of follow-up, and there's some good discussion about that. I think the key issue on semiconductor chips is obviously, and this is, again, the overlap between globalization and technological advance. Do we care where they're made? If a lot of them are made in Taiwan and we're worried about China acting against Taiwan or whatever, or earthquakes in Taiwan, how much of a problem is that? What kind of disruption would that be? How much of a switch would that force upon our economy? You know, I think the really good news is that a lot of global companies that are very good at making things like Taiwan Semiconductor are actually, okay, they might complain a little bit, but they're actually quite keen to have operations in the United States as a form of portfolio diversification themselves. There's a lot of European companies piling into the US making in the clean energy manufacturing space, which doesn't make the European governments very happy. But as far as I'm concerned, it's generating jobs in the United States and there's a potential race to the top there in terms of encouraging other investments. cheaper technology and pushing that technology space in the right direction. So I think the semiconductor industry, in part because of what happened during COVID, in part because there were a lot of cars that were manufactured but couldn't be delivered because they were missing a couple of chips. I think they won the argument that that's a strategic potential bottleneck. And then the question is, how much do you want to spend? What's the net cost? to having more of that production, particularly if it's advanced chips in the United States, if you're willing to subcontract out the more generic chips to other places. And does Taiwan in particular? How do you feel about the geopolitical risks there? If Taiwan were located somewhere else, if Taiwan were in a relatively calmer part of the world, I think having more production there and having them specialize in that would be totally fine. Good ally. I mean, Ireland, for example, makes quite a lot of ventilators, which a lot of people will climb with ventilators during COVID. As far as I know, the Irish never... I mean, these were the subsidiaries of American companies, but there was never an issue of not delivering ventilators of

contract to the United States, even when the Irish really wanted it, the Europeans really wanted it. So having allies with whom we have a deep, collaborative economic and political relationship, and where the sanctity of contract holds up just like in the United States, that's what you want. I think the problem with Taiwan is just the neighborhood, right? Yeah. I would just point that we're not entirely sure how the labor market is going to respond. And the fact that if you react too soon, the prevention programs may be an unnecessary event that's not done properly.

Audience Member | 49:00.000

So given the limited resources that the fiscal environment can find ourselves in, what is an appropriate investment in the education program to prepare the future workforce? Is there anything we can do right now?

Robin Hanson | 49:20.525

work study. I mean, you know, basically people have to be trying this out and until we know how it plays out to actually use it, we won't know how to teach it. So just help students get general skills. So this is something you mentioned before. I used to be in computer science a long time ago and I would see people get into computer science and spend all this time learning very particular computer skills that were context dependent to the particular economy they were in. And I was kind of ashamed because even though they were learning computers, they were learning particular skills that wouldn't necessarily last. So I would suggest in general, in times of uncertainty, invest in generic skills, general skills that will be useful in a wide range of contexts. Math, write, reading, writing. Present, you know, communication, teamwork, those are the things that you can be pretty sure will be valuable no matter where the technology goes in particular, when there's huge technological uncertainty is exactly the wrong time to be making bets about which particular technologies will win.

Simon Johnson | 50:21.394

So the DARPA, the first DARPA self-driving car competition, I think had a prize money of about a million dollars. One million dollars, right? And when the people who were on the MIT team, which didn't win, when they put up a slide like talking about their lives since then, there was a picture of like 30 people on the photograph and 14 of them have set up their own companies subsequently, that sort of thing, right? And I think there are moments when you can catalyze things with attention. And with a little bit of prize money, sure, but it wasn't about the prize money, right? It was about the prestige and it was about like opening this area and saying, hey, this is coming. Actually, I don't think we were ready for self-driving cars, but that competition helped move us towards being closer to it, perhaps. So I think I would do that for education. Education and health care are two spaces that are potentially where you could define competitions and say, look,

we'd like to help people learn this. We'd like to help lifetime learners. We'd like to help people pick up. These are the following skills or maybe it's just broader than that. And I think getting people's attention and focusing on that and emphasizing, if you want to emphasize the national security angle, that's totally fine with me too. I mean, DARPA was going with self-driving cars. I mean, they were not. Military vehicles per se, they were cars that could drive around, okay, in some military type situations. But we'd like our soldiers to be able to learn things fast. We'd like the military to be more adaptable in this kind of situation. We have a huge number of people in the military and we want them to be better prepared when they go to the Civilian Labor Force after the military. So it's not crazy if you think the DoD and DARPA is a way to do this, great. If you'd rather go through St. Mary Avenue, totally fine. But I think getting public attention and catalyzing creative people and raising the legitimacy or credibility of an issue is a very good thing for the government to do.

Jason Green-Lowe | 52:08.433

In the back.

Audience Member | 52:11.265

So, curious maybe for some rough forecast of different governments like we have to do. I think for context, the US is a very tight labor market right now, low new employment, China is a pretty fine one, people can pay, fill up the lease. I think there's a lot of ways much different countries are exposed to automation. Just curious for like, how do you see different countries responding differently to AI if it ends up automating?

Robin Hanson | 52:45.076

That if is the key question. I mean, the key question at the moment is, will this get traction to the degree that people are hoping or not? If it ends up being a bubble, then whoever invested a lot will have wasted a bunch of money and time. But if it ends up getting a lot of traction, then people will maybe kick themselves for not being in sooner. That's the most fundamental thing is different places in the world are making different bets about how much to throw in on this. You know, U.S. is clearly making a big bet, China to a substantial degree, but... India less so. India, I might say, is betting, yeah, this won't pan out. Let's not waste our money on this. And we don't know yet who will win.

Simon Johnson | 53:23.994

Yeah, it's a good question. I think every country is going to be different. I think one interesting dimension is the European Union, because obviously they've adopted a more restrictive... general tone, but I think they also have a really interesting potential here because the extent that workers, when you say to people like, congratulations,

we've got an AI strategy, we're bringing AI to the company, I think a lot of American workers initially are actually going to be, oh, right, there goes my job, right? So you've got a little bit of a trust gap. In European situations, so the German Workers' Council would be a good example, but there's other examples. If the workers are more convinced that the technology could be used to boost their productivity and their wages, and they will not necessarily get... lose those good jobs, then they may actually be more willing to adopt some of these technologies. So there are people in Europe who are making that argument, and we'll see if that prevails. Now, of course, Robin's right. Is this a big deal? Does it help a lot? And so on. I think over time, it does become significant, and it's worth pursuing. And it could be a place where Europe actually gets a bit of an edge, which would not necessarily be so obvious, given that we are the ones who have 90, 95% of the technological... develop the capacity to develop the technology right now.

Audience Member | 54:35.000

potential arms of jobs as a person, in the future of automation. We've each spoken a bit about the need to create new jobs, new time assists, new potential for displacement. We've talked much about universal basic... I'm wondering how does the UBI compare to job creation or combination of other interventions?

Robin Hanson | 55:10.731

I focus more on enabling job mobility, which includes geographic mobility. That is often in these new bursts of new technologies, the benefits are unevenly distributed in different places. Some places are boom towns, some don't gain as much. And the boom towns need people to move in and... At the moment we have a serious housing problem in the US that we find it hard to build new housing in the places where a lot of new jobs show up. If you were going to enable us to adapt more rapidly to the changing economy, I might say what else besides jobs? I might say housing. Find a way to make it easier to build more housing where the people want to move.

Simon Johnson | 55:50.831

Yeah, so you, Groupon Housing, let me come to that in a moment, but on UBI, I'm not a fan. And the reason is several fold. First of all, I think people want to work. There's dignity and identity that comes from working. Secondly, the productivity gains you're going to get in the foreseeable future are not that big. So there's no like big, massive amount of additional income to hand out to people. And I think thirdly, the politics of a UBI, where you have a set of people who have jobs and a set of people who don't. And you say, well, these people don't have jobs. It's not really their fault. We'll keep them. well paid or well remunerated indefinitely, that's not going to work in the United States. I really don't. I mean, maybe in some other parts of the world, you might imagine that

works. I think politically it's insupportable. So I think what you want to do is focus on new tasks, new task creation, new tasks with expertise. That's what's going to help people. And including these other things we've been talking about, like, you know, lowering the barriers to entry into those kind of activities. I think housing is a first order problem. And I'm glad that Robin's brought it up because what used to happen in the United States, What happened for 200 years was every time there was a boom with a technology that was associated with a place, Chicago for example, Houston, Los Angeles, each of these places were nothing and then there was suddenly a hot place and it was linked to a technological transformation. Millions of people moved there and we built housing. What's happened in the past 40 years is we've had some booming technology on the east and west coast, but those booming places like Boston or New York or, worst of all, San Francisco, they have not built housing. So as a result, in the 1960s, if you were a... a low-income cleaner and you moved from the deep south to California your income post housing costs absolutely went up but that's not been true in recent decades because of the cost of housing now the problem is it's these are most all local restrictions it's almost all zoning restrictions and it's local politics and it's the people already incumbent and including the homeowners with whatever attitudes they have to whatever it is who don't want to add higher density housing So I do think that you can give more incentive for high-density housing. That's very sensible. I think you're struggling, and you can see some local mayors, New York, for example, is attempting to deal with this. Some others are trying to move somewhat. But the population of the United States increased by 2.5 million people last year. These are small drops. What you really need is to have more of this tech expansion happen in other places. I spoke up for Huntsville, Alabama in my Senate testimony. That went down very well with the Senate from Alabama, as you can imagine. And we wrote a book about this called Jumpstarting America, which came out in 2019, arguing that there are... 35 states and more than 100 places that are plausible, you know, incipient new tech hubs. The problem is talent wants to go where talent already is, right? So you're going to get some crowding into the existing tech hubs. But to the extent that you're making government strategic investments, to the extent that you're supporting universities around the country that are good already and that have ability to develop further, I think encouraging the development of more geographies is very smart. Because the market... wants everyone to go to San Francisco and Boston, but the political market on housing is very bad. So, looking for ways, I mean this is a country that, unlike where I'm from in the UK, which has got a couple of big cities that have always been dominant, we've never had that kind of model for our economy. Actually, it's relatively recent. We've actually moved the center of our economic geography around a lot until recent decades when it became rather too rigid.

Jason Green-Lowe | 59:29.945

All right, so that's the official time we have. I want to thank our guests for joining and educating us today. Thank you very much for coming. The Center for AI Policy is really pleased that we were able to spark this discussion, and we'll be here in the coming months keeping the conversation going.