

THE BROOKINGS INSTITUTION

WEBINAR

WORK IN THE AGE OF ARTIFICIAL INTELLIGENCE

THURSDAY, APRIL 25, 2024

MODERATOR: ANTON KORINEK

Nonresident Fellow, Center on Regulation and Markets, Brookings

PANELISTS:

DARON ACEMOGLU

Institute Professor, Department of Economics, Massachusetts Institute of Technology

DANIEL SUSSKIND

Senior Research Associate, Institute for Ethics in AI, Oxford University, Research Professor, King's College London

RUSSELL C. WALD

Deputy Director, Stanford Institute for Human-Centered Artificial Intelligence (HAI)

* * * * *

KORINEK: Hello, I'm Anton Korinek. I'm a nonresident fellow at the Center on Regulation and Markets at Brookings, as well as a Professor of economics at the University of Virginia. It is my pleasure to welcome all of you today to our event on work in the age of artificial intelligence. And we are indeed living at the cusp of the age of artificial intelligence. AI is everywhere these days. In the news, in the workplace, and every week we see new AI models announced. So to discuss this question. We are honored to have three distinguished guests with us today Daron Acemoglu from MIT, Daniel Susskind from Oxford, and Russell Wald from Stanford. I will introduce each of them in turn before asking our speakers for five minutes of initial remarks each. And incidentally, if you are watching this event live, you can tweet any questions that you have under the #Alandlabor, and I will do my best to incorporate them in our conversation. So to start, I would like to invite Russell to kick off our conversation since he has just published the annual AI index report at Stanford last week. And I thought that the findings from this report may provide a really excellent grounding for our discussion. I have followed the AI index report for many years now, and it is one of the most authoritative sources on what's happening in AI, providing an annual snapshot of the state of the field. And you can also see the growth of AI in the growth of the index report. The 2024 edition came in at 502 pages. So Russell Wald is the director of policy for Stanford's famous Institute for Human-Centered Artificial Intelligence, HAI. In this role, he is responsible for leading the team that advances HAI's engagement with governments and civil society organizations to see a world benefit from the human centered uses of AI. Russell is also a member of the steering committee of the 2024 AI Index report. So, Russell, could you share a few highlights from the AI index report with us and tell us a bit about the speed of AI advances and the implications for work in the age of AI?

WALD: Thank you. Thank you, Anton. And it's a privilege to be here with everyone and such a distinguished group. I will say that I am now the, deputy director of HAI, which was a recent development. And, so we're looking for a director of policy replacement, and this is probably, audience that would be great for that. So if anyone's interested, find us. In the -- but, yes, I do work for Stanford. HAI Stanford's Institute for Human-Centered Artificial Intelligence. And with that, we do help publish the AI index, which, as you noted, is this preeminent measurement of artificial intelligence. It is indeed 500, over 500 pages this time, so it's quite lengthy. It's reflective of the growth of AI. I'm going to give just a few points that I see that came from the report that are quite fascinating right now. First, one thing that we're seeing is, AI in numerous fields is exceeding human performance at this point. We saw this of quite a few years ago with ImageNet. And, at that point, there was point as to which human performance was eclipsed. And now we're seeing this from, a variety of different areas, whether that's, reading comprehension, and ultimately at some point, we're getting

closer and closing the gaps on a whole nother set. What's interesting is, is at the same time, we also see an increasing difficulty of benchmarking these systems. And how are we able to appropriately understand what the level baseline of benchmarking is? That's important because one area that we know in the report is there's no one standard universal level benchmarking right now. And why that's an issue is there's a variety of these large firms that are developing this technology, but they're not doing it across the standard level of benchmarking. So we don't know what the performance capabilities are of all of these across one standard. And that will become a growing challenge in the future. And it also begs the question of who should be creating that one universal standard for us to rely upon? Is that government, is that industry, is that civil society in some capacity? That's, another key feature that came from this year's report. Then I think what we should talk about are foundation models. Foundation models are large, very large models -- think ChatGPT today, kind of, for a more general audience here. And what it is, is they are use an extraordinary, extraordinarily large amount of computational resources combined with a huge corpus of data, generally from the internet. And they are trained at very high levels. And later they're intended for downstream task, to do a multitude of downstream task. And so we've entered an era in a world of multimodality, meaning that there is not just a bunch of different models, but over time, you'll see one model be able to do many different things. And we're seeing an incredibly huge uptick and rise in that. And we're also seeing an incredibly huge uptick and rise in that related to the cost of these. And that gets us to my final point here. The cost of these models has run away well beyond academia. And industry clearly has the overall dominant lead here in this space. It sets up a variety of policy questions which I hope we can get into in much more detail to talk about the need for public sector AI investment in this space and why that's so critical. But just to give you an anecdotal piece from this, what we note is in 2017, Google spent \$930 to train their transformer model. Well, the transformer model is the T in GPT. It is the underlying base of, underlying level of technology for generative AI and heavily used, quite often right now. So \$930 is what it took in 2017. Fast forward five years later and what do we get? We see that Gemini Ultra from Google cost nearly \$200 million to train. So we've gone from in five years, \$930 to \$200 million to train some of these in the industry level. This has huge implications and significant challenges for, a variety of reasons. And I do hope we can have a more rich conversation about that.

KORINEK: Thank you, Russell, for the fascinating remarks. Let me now turn to Daron, who will help us shed light on the economic implications of what we have just heard about. Daron is the Elizabeth and James Killian Professor of Economics and Institute professor at the Massachusetts Institute of Technology. He was awarded the John Bates Clark Medal in 2005. That's a greater honor than the Nobel Prize in economics,

because it is scarcer -- and economists measure everything by scarcity -- and by some measures of citation counts, Daron is the most influential economist in the world. Daron has worked in many fields of economics, but his research on the effects of AI on labor market has been a particularly important inspiration to me. He has operationalized the so-called task-based model of jobs to analyze the implications of automation and AI for labor markets. Daron, would you mind sharing your concerns about the effects of AI on the workforce with us?

ACEMOGLU: Thank you, Anton. Well, I think Russell really set the stage brilliantly. The amount of investment and the rate of advances in artificial intelligence in general and generative AI in particular are spellbinding. They're truly impressive advances that even insiders would have thought impossible 5 or 6 years ago. And there are some truly promising directions this technology can go. But with any technology, its effects are going to depend on how we use them. And when it comes to production, for general purpose technologies or technological platforms, there are many choices about how we can operationalize them. In particular, we can use technologies for automating tasks, meaning that we take tasks that were previously performed by workers and we let machines or algorithms take them over, or we can use them to provide better information or greater capabilities to workers that are performing some of these tasks. Or we can do something even more important for the economy. We can create new products and new tasks. So if you take the internet, for example, I think the internet took some time to have an impact on the economy, but it was most transformative because it enabled a range of new tasks and new products and a lot of information for workers. We've used the internet for some bad purposes as well, but on the whole, I think the internet's brilliance was its flexibility in allowing us to do many new things. So if we can use AI in the same way, I think it could be a positive development both for firms and for workers. But -- here's a big but -- I am very concerned, because that's not what I see being done at the moment. I think a lot of AI research is spearheaded by two considerations. One is artificial general intelligence, the sort of prioritization of making these models humanlike. And second, large businesses wanting to automate and standardize many aspects of their jobs in order to cut labor costs. Both of those are visible, for example, when you look at ChatGPT, you know, the, you know, there are a lot of design choices there. And those design choices were often made in making these foundation models, and the applications built on them, sound very human like, and, you know, try to pass a Turing test like contraption. And the first ways of monetizing them was either for automation or using them for digital ads or related business models. These are not going to be very good for workers. Automation, of course, has been with us and will be with us, but it benefits capital at the expense of labor, and often it creates even large within-group inequalities, because the workers who are displaced from

the tasks for which they have a competitive advantage end up working for less or lose their jobs completely. So it is really essential for building shared prosperity on the basis of AI, that we abandon the obsession with artificial general intelligence and automation and try to use them for improving information for workers, creating new tasks, helping the workers, empowering them rather than disempowering them. So that's why I think the question of what will AI do to the labor market is an underdetermined question. It will do whatever we choose it to do. And I think there are very important choices that we should make as a democratic society and not just a few leaders in Silicon Valley. And I'm really worried that despite all of the discussion about AI, that basic truth that those different choices cannot be just delegated to a few people has not sunk in. And I think that's a very important part of the whole discussion about social media, democracy. But it's also the heart of the inequality and jobs discussion.

KORINEK: Thank you, Daron. Let me now turn to Daniel Susskind. Daniel explores the impact of technology, and particularly AI, on work and society. He is a research professor of economics at King's College London, a senior research associate at the Institute for Ethics in AI at Oxford, and an associate member of the economics department at Oxford. I have followed Daniel's work on the future of work for a long time, and I particularly enjoyed his book entitled "A World Without Work," in which he cogently lays out how the coming wave of automation may make labor redundant eventually. So for a glimpse at the future, I thought there are few better speakers to invite than Daniel. So, Daniel, we've just heard from Daron about his short-term concerns about AI. You have extensively written on the longer-term future of work, and I think you have a very a new book on economic growth coming out as well, if I understand correctly. Often times, people who worry that AI will displace work are accused of being Luddites and not understanding basic economics. Now, this is certainly not true of Daron, and this is certainly not true of you, but can you explain why you view a world without work, or with much less work, a plausible outcome, and what economic forces you see that may lead us there?

Susskind Thank thank you very much for the warm introduction. And it's a great pleasure to be with everyone this evening from London. I, in the short term, my, my view is that the, the challenge we face is not a world without enough work for people to do. The challenges that there is, there will be work for people to do, but for various reasons, that work might sit out of their reach. That because of the sorts of work that technology creates, there might be some, what we could call frictional technological unemployment. I mean, the most obvious reason is that people might not have the right skills and capabilities to do the work that has to be done. But I think there are other mismatches too that are, that are important. Another that work is

created, but just not in the particular geographical place, that people happen to live. And then I think there's also really interesting identity mismatches as well, where people have particular conceptions of themselves and they're willing to stay out of work to protect that identity, given the sorts of work that's available in the labor market. But but I do think that as you look further into the 21st century, the challenge might change from being one where there is work, but for various reasons, people might not be able to do it, to one where there just might not be enough demand for the work of human beings, full stop. And and in in the book that you're, referring to, "A World Without Work," I set out what I think are the big challenges that we would face if we found ourselves in that sort of world. One is the economic challenge. And it's a challenge fundamentally of inequality. You know, I, I don't think it's a coincidence that today worries about inequality are intensifying at exactly the same time that these worries about automation are growing. There are these two problems, inequality and technological unemployment, are very closely related. Today, the labor market is the main way that we share our income in society. For most people, their job is their main, if not their only, source of income. I think the inequalities we see around us today show that that approach is already creaking, that some people get far more for their efforts than others. Technological unemployment, in my view, is just a more extreme version of that same story, but one that ends where where some people get nothing at all. So I see the economic challenge of, you know, technological unemployment as fundamentally a distributional challenge, a challenge of how we might share out prosperity in society if the labor market does a less effective job. But I also think there are other challenges that have less to do with economics. One is the challenge of power, and in particular, and it's been referred to already the sort of political power of the large technology companies who are responsible for developing these technologies in the first place. I think in the 20th century, our main preoccupation was with the economic power of large companies. Now we worried about excessive profitability, you know, market concentration, predatory pricing, and we develop tools and, you know, legal frameworks for thinking about how to identify those concentrations of economic power and intervene appropriately if need be. I think what's interesting and challenging about the 21st century is what's far more concerning about these large technology companies is not their economic power, but their political power and the impact they have on things like liberty and democracy and social justice and whether or not those things are under threat. And I don't think yet we have the same toolkit for thinking about those uses and abuses of political power that we developed in the 20th century for thinking about economic power. The final challenge, and again, this has not too much to do with economics and it's a challenge of meaning and purpose. Now, it's often said that, you know, work isn't simply a source of income, but it's also a source of direction and fulfillment. And and if that is right, then the challenge isn't simply that technology might hollow out the labor market, leaving some people without an income, but it might also hollow out that sense of

meaning and purpose, too. And I think buried within, the discussions that many people have today about the sorts of interventions that might be required in a world with less work, things like universal basic incomes, things like job guarantee schemes, are often tacit assumptions about the nature of that relationship between work and meaning, which I think are really important to set out and and interrogate.

KORINEK: Thank you, Daniel. And thank you all for your initial remarks. Now to start the conversation part of our event, I thought I want to ask you all a quick and somewhat lighter question that focuses on our personal experience. How much have you personally used these recent advances in generative AI in your own work and life? And what lessons do you draw for your research? And I should say that I personally, I have spent really quite a significant amount of time exploring and employing these generative AI tools like Claude and GPT-4 to make my work more productive. I have published a paper on the topic, and it's actually also listed on my Brookings page, and I estimate that I'm at least 15 or 20% more productive using these systems. So let's perhaps proceed in the same order in which we had our initial remarks. And let's try to make this a quick round, of about a minute each. Russell.

WALD: Sure. So I use it sparingly. And I use it sparingly more because of the hallucination factor on large language models and the confidence that I have. And so if I'm using it, it's on something that's short and that I can easily fact check and edit myself. If it's extensive, then, I don't rely on it too much on an active way. You're seeing actually, in, you know, for example, Google Docs, there's actually, they're getting much it's getting more sophisticated at finishing the sentence for you and stuff like that. But if I'm actively going into Claude or ChatGPT and using them, I'm very careful as to how I use them because of the hallucination factor, and I don't want it to come back that I got something wrong. That doesn't mean people shouldn't use them. I think they should learn how to use them and not run away from it, but be mindful of its limitations.

KORINEK: Daron

ACEMOGLU: My experience is very similar to Russell's. I used it, you know, more than a year ago when ChatGPT-4 just came out and, you know, I was at first very impressed, but the unreliability of the information put me off and I'm not using it at the moment. But I think, like many people in the US economy, I'm not using it, but I'm subject to it because many of the programs that we use, websites are all using AI algorithms. If you look in the data, it's like less than 2% of US firms are investing in AI, but a very large fraction of US GDP and US population are subject to AI in a way that they don't understand. They can't monitor and they don't have a

say on it. So I think that's part of the problem that I tried to highlight in my previous answer, that right now, AI is a tool in the hands of a very small segment of society that can that can use it in unaccountable ways. They may choose to use it for good ways for most of the time, but we have no recourse against it.

KORINEK: Yeah. Daniel.

Susskind I think what I found interesting is I found myself using it quite often for tasks that would normally require creativity from me. So I've used it for, I used it when it first came out to help me draw up a long list of titles for my new book. I thought I thought that was quite, you know, helpful and, you know, pushing it and saying, you know, I want something more, you know, engaging, something a little more intellectual, something a little more like, you're just, you know, pushing it and getting a long list of titles. I mean, I also I've used it and I don't know if this makes me, a good or bad parent, but I've used it to help generate bedtime stories based on the prompts that my, you know, six year old gives me. I think the serious point here is that what's been quite provocative about GPT is that it's particularly good at some of these creative tasks. And if you had gone back five, you know, even a few years, one of the areas, one of the human faculties that many people would have thought was out of reach of even the most capable technologies for some time were those that required some kind of creative faculty. You know, solving the problem of originality or novelty, or taking one by surprise. Turns out that actually, you know, these systems, whether it's generating text or images or videos, are particularly good at some of those, some of those, activities.

ACEMOGLU: Well, if I can share just one anecdote, I was about to submit a paper to a journal where it would requires the abstract to be 100 words. So I emailed my coauthors and I said, you know, in LaTeX, it's very difficult to check that. Have you checked how many words it is? And they said, yes, yes, it's 95 words and I upload it and the system says, this is 140 words. And I said, I write back to them. I said, you said this was 95. I said, oh, we checked it in ChatGPT, it said it was 95 words. So, you know, there are things that you can make mistakes, but the degree of certainty and the mismatch between what is sometimes true and what's not is still striking.

WALD: I know this was supposed to be quick, but I want to add one more point to build on what Daniel said about using it. That is one way, actually, I have used it as well, but not for an outright finished product, but more as an assisted tutor of helping me with creativity. And we have this, and I have an agent here that I've gone back and forth with. And through that, it's unlocked things in my mind that I would have never thought

about and said, oh, that, oh, we could go in that direction. And so it is kind of like having, you can whiteboard by yourself, but it's a whole lot more powerful if you whiteboard with someone else in the room.

KORINEK: Yeah, that makes sense. And Daniel, I should say, my children also love generated bedtime stories, and I make sure to prompt the GPT to also add some morality to them. But now let me turn a little bit more into labor market effects and, to more near-term questions. So, if we look at the next two years or so, how do you think, can we steer the labor market effects of AI in a desirable direction? And I should emphasize, both in the economic domain and more broadly, when it comes to non-monetary factors that affect workers. Let's start with Russell again and go in the same order as before.

WALD: Oh, I hate to talk about labor markets with two brilliant economists here. But what I would say is, what we're seeing that's very fascinating in early results of this is that, it's not so much a replacement activity here, it's much more of an augmentation tool. And that's what there is some hypothesis around that. But you're seeing, in some of the early studies when paired -- and it's noted in our AI index -- when paired with the worker, that you see a greater degree of productivity, you see a greater degree of efficiency, speed, quality, that's coming from this. And one example was an experiment that was done, and that's noted in the index, is there was a subset of a group that was called, that, to use an AI system. And it was not, to, to not use an AI system to work on something, another group to use the AI system. But they were told it has issues. Be careful. And then there was another group that was told, use an AI system, it's quite, it works extraordinarily well. And what turns out is, is the ones that had the AI systems outperformed those that didn't. The one that actually used the AI system that was told it's a little shaky, so be careful, outperformed both groups substantially because they were much more careful how they did it. So that's possibly where we see a lot of this going right now in the immediacy of a more of an assisted tool than an outright replacement.

KORINEK: Daron, do you also see the near term to be augmentation? What do you think are the implications for the skill premium and so on? And what do you think, do we need otherwise to make sure that we steer this technology in a worker-friendly direction.

ACEMOGLU: Yeah. I think, you know, we're very impressed with this technology. And there's a tremendous amount of investment going, but I expect it to be a much more slowly changing phenomenon. So I don't think we're going to see massive effects within two years. I suspect there will be intensification of use of AI for surveillance in workplaces as well, and some creative workers such as yourselves and the sort of very

sophisticated employees are already using it for a few things. I think, you are already very productive, Anton, so 20% increase in productivity is too much. So I'll give you 5% perhaps. So many people might have 5% improvements in their overall productivity. But it's they're they're they're a small minority and they are not the workers at the bottom of the distribution who we are most worried about. I think the effects for them is going to be much more of a slow burn. And there that's both good and bad. It's bad because it's a promising technology. It's going to be slow. Fine. But I think it's actually good because it gives us more time to adapt. And I think the augmentation problem, Anton, I think that is key. And right now, we don't have a road plan for that. We don't have a strategy. Tech companies are not prioritizing it. I can talk about in detail across different occupations and sectors lots of low hanging fruit. And nobody is sort of going after them. So the slow pace perhaps can allow new companies to enter and try to see whether they can use generative AI to make teachers more productive, to make nurses better at dealing with shortages of health services in the United States and many other countries. Courts that have huge delays become better, electricians become much more productive in dealing in problem solving tasks involving new technology. Plumbers, carpenters. I think there are a lot of possible applications, but those are not where the money is going. So yeah, I think within two years, we'll probably be where we are now bemoaning the fact that they're not doing enough augmentation.

KORINEK: What specific policies would you advocate, Daron, that would --

ACEMOGLU: Well, that's a longer discussion. I'm more than happy to get into it. I don't think there's a silver bullet. So I think we need to deal with industry structure. We need to deal with fiscal incentives that are already in the system. We need to deal with business models that are very distorted in the tech sector, financing models like venture capital. We need to have the government play more of a regulatory role. And also the government played more of a role like it does in climate change, energy technologies, or health care by providing subsidies, prizes and other incentives for socially beneficial use of AI. But I think the first step, and I think this is a really, really big step, is if there is a broad agreement in society that it is both socially desirable and technically feasible to use AI and generative AI in particular in a pro-worker way. I think we've we've made huge progress. I think that reality is denied, that reality is denied in many different ways. There are those in the tech world who wouldn't even accept that it's socially desirable to invest, to make, in order to make these technologies more usable by workers. Those that would say that's not technically feasible, we just have to go down our foundation models or those who would say, oh yeah, everything's gonna work out, just let it, let it remain in our hands and we'll just do whatever is best for society. Good luck. So I think if we

can overcome, these sort of, roadblocks, I think we've already made a big, leap towards more pro-human pro-worker AI.

KORINEK: Thank you, Daron. Daniel, what is your take on that? And you know, I can't help but I have to ask, what is the title of the book that you finally settled on? And do you want to give us a just brief plug on that new book as well?

Susskind Yeah. So the book's called, "Growth: A History and a Reckoning." And, you know, one way to think about that book and, you know, my, my starting point in this discussion about directing technical changes, that I'm very sympathetic to the idea that we ought to try to shape the direction of technological progress. You know, one way to think about the argument of the book is that, you know, I'm saying we ought not to simply shape it with respect to the the impact of technology on the labor market or on politics, but also technology with respect to its impact on climate, on its impact on local places and communities. You know, all the all the different harms that we might think that technological progress has, you know, it has this extraordinary promise, but it's also got this price as well. So I'm very sympathetic to the idea that we ought to be trying to, you know, direct technological progress in, in particular directions. My I suppose what worries me is, and it relates to what some of the comments that have been made, it's really our capacity to do it. In part it's the sort of the practical limits. It's quite interesting when you, you know, in, in economics, we draw quite a clean distinction between the idea that technologies can either, a complement for human beings, making them more complement human workers, making them more productive, the work that they do or substitute for human beings, displacing them for particular tasks and activities. What's been quite interesting talking to computer scientists is that they sort of that they find that the cleanliness and the sort of fixed nature of that distinction quite misleading. I think it was the computer scientist Stuart Russell in his BBC Reith Lectures who made the point that it's just it's often in practice very difficult to know ex-ante whether a particular technology is going to end up being, you know, being a substitute or a complement. You know, you might think of, for instance, a, you know, a medical diagnostic system which you were originally designed to help, you know, help a doctor, you know, make a better diagnosis of whether a freckle is cancerous. But it ends up doing such a good job that, you know, the human doctor gets in the way. And what was meant to be a complement ends up being a substitute, and patients want to just use the technology instead. So, I think there's there's been, whether we can really have the kind of clarity on whether technology, you know, what, and knowing for certain whether technology's going to have particular effects, I think, is, you know, one of the one of the things that, you know, I'm grappling with, the other is the politics of it. You know, and I think there

are interesting, you know, even if it were possible in practice to direct technical change, is it politically feasible to do so? And it's interesting to look at, yeah, because we could tell exactly the same story about trying to direct technical change with respect to the impact of technology on the climate. You know, for decades we've been able to direct technical change and, you know, try and encourage greener rather than dirty technologies. And we've made an extraordinary amount of progress, but, you know, many people today would argue we've not done enough, and we should have done more over the last few decades. And yeah, you might argue we're a similar moment now with respect to AI that we were with respect to climate change back in the 80s. And, you know, the sort of political impediments to directing technical change with respect to the climate, do those not also imply, also, apply with respect to the challenge of technology in the labor market? So, so I, you know, in the short to medium run, I'm, as I said, very sympathetic to the idea. I worry a little about the kind of, you know, whether it's, and, you know, I advocate for the idea, but I worry about the sort of practical limitations and the political limitations.

ACEMOGLU: So can I just jump in on that briefly? Because, you know, I think Daniel is obviously raising very important points, and there are many economists and commentators that would much more strongly caution against trying to steer or redirect technological change because of political economy reasons, inefficiency reasons, or it's just impossible. Yeah, obviously, I think there are many, many difficulties and there will be some degree of unpredictability, but there are very good guidelines that we have that we can already start thinking about and we can improve. Like, you know, in education. I don't think it's going to be a surprise that if you do more and more automated teaching, that's not going to be good for teachers, and it's probably not going to be the most productive thing for students. And there are many opportunities for developing AI tools to help teachers become better at identifying student difficulties with exactly which part of the curriculum, do much more personalized teaching. And how much money is going into that right now? Completely trivial relative to automated teaching expenditures. So you see the market, you see the mindset, and I think there are obvious things to support. Exactly like in the climate change. We haven't done enough in climate change. We've made mistakes. Some money went to waste. But if we hadn't done that redirection today, we would be looking at 3%, sorry, three degrees relative to pre-industrial times, not still debating whether we can do 1.5. So the, the success, limited though it may be, was all due to our willingness to invest in solar, wind, geothermal and other renewables and battery technologies. And I think this is the same situation is true today. We'll make mistakes, but I'll welcome those mistakes.

KORINEK: Thank you Daron. So, I want to bring in a few questions from our listeners now. And I'll start with two technical questions. And I'll direct them at Russell, since you have just produced this authoritative overview of the state of AI. The first one, which came in, quite recently, you mentioned before that estimates for the cost of training the most expensive AI systems to date are around 200 million. I've seen estimates up to 500 million for that same system, actually. And one question is, how much of that is capital costs versus labor costs? In other words, who benefits from training GAI systems at this stage of spending money on the training? And then a more broader question, perhaps, which industries and job roles do you see will be the most impacted by AI in the near term in the next couple of years?

WALD: So, this is, reflective of the training cost, not the labor cost or, the other wider parts of this. So that alone tells you that it's much more significant.

KORINEK: Purely the cost of the computer.

WALD: The cost of the compute parts of this. But what is interesting and remains a consistent challenge for us is a lack of transparency in some of these cases from some of the companies. And that lack of transparency makes it difficult for us to measure what's happening in this space. We'd love to know how many GPUs ChatGPT will be, assuming GPT-5 will be trained on. It will be really important. It'll tell us a lot about the energy consumption. It'll tell us a lot about, what's going into, even Meta that will talk about and celebrate its openness, it's open so far as the model weights are open, which is just essentially a part of the architecture of this. It is not open in a sense of what that was, what went into these systems to train it. So what this bigger part question is highlighting for me is a lack of transparency that we are consistently struggling with. And what can we do to be able to ensure that industry is becoming more transparent in order to better vet these systems and audit them and be, have different stakeholders at the table participating to understand what happens here. And related to the second part of the question. What do I see as the most, promising or, areas, parts of AI? First, I'll say that I think we have a very strong, extraordinarily significant obsession, particularly by policymakers, that I think are hearing exclusively more from industry, and that is this, captivation towards the language model. There's so many other parts of AI right now that I think we are not paying enough attention to, that are going to kind of once again, surprise us if we're not aware of that. So not just the multi-modal foundation models that I talked about earlier, but also areas of ambient technology understanding your, your physical space and what that looks like through world modeling and what can come there. There's a whole different world of AI that we're not spending enough time thinking about, that we

should be becoming more aware of. The, and then the to add on to that second question, what I find the most interesting is actually the, AI-driven or AI innovations that are happening in science that actually had AI behind it. We have actually an entire chapter in this report dedicated to science and medicine and what's happening here. But if you look at, for example, the nuclear fusion breakthrough that happened last year, we sit here and we refer to that as a nuclear fusion breakthrough. And this is incredible thing that happened. But that likely would have never have happened at this point if you didn't have AI and ML techniques applied. So what I think we're probably going to see is a lot of this applied aspects of AI innovation in different fields and domains. And we won't sit here and talk about the model that was used or how it was used. Instead, we're just going to talk about a unique breakthrough, scientific breakthrough that happened. In essence, you're unlocking a lot of new knowledge as paired with this technology, just as we were talking about earlier, and Daniel noted the creativity that can come from this just as me, as a user, the same thing can happen for a scientist that's working on it in a different way, and use it to understand their data in a different way, provide different insights and help create more breakthroughs. So what I expect to see, and we have seen early data on this is scientific discovery is about to get supercharged in a way in a variety of fields. And we're not going to sit there and say it's AI, we're just going to sit there and most likely note or comment that it came, this unique scientific discovery happened. So I think we're being a little sometimes too myopic on the language model. And we need to open up our eyes wider to the other applications that are going to happen here, not just from adoption of a language model.

KORINEK: Thank you, Russell. Now, a follow up question to Daron. I hear this quite a bit that we will see massive effects on the process of scientific discovery and also innovation in the entrepreneurial realm because scientific breakthroughs need to be propagated into the economy by innovators who actually implement them so that, we as a society more broadly can benefit. Now, Daron, as one of the preeminent growth scholars, what is your reaction to hearing that?

ACEMOGLU: Well, it would be fantastic if it happened, but I wouldn't bank on it. First of all, you know, I see, at least at this current level of AI knowledge, zero probability that AI could make a scientific breakthrough. What it can do is it can narrow down the field of search, but you still need very skilled people to do the experimentation and, you know, work on the crystal structure stress suggested by AI, or the particular protein foldings that are feasible. So the human element, the labs, etc., all of those are going to be necessary. So the gains would be incremental. Second, you know, the last five decades have been one of huge disappointment on the link between innovation and productivity. Almost all measures of innovation have

skyrocketed in the USPTO, United States Patent and Trademark Office. There are five times as many patents today as there was in 1980, and the rate of aggregate total factor productivity growth is as slow as it has ever been in the 20th century. So, so we're somehow not able to get what we care out of these innovations. So I think a top-down innovation process controlled by a few firms is unlikely to change that picture. So again, I think the promise is there. If AI becomes a decentralized tools that enables new entrants, new firms, new products, new perspectives to unleash more energy, yes, absolutely, I think it would have an impact. But that's not the path we're going to.

KORINEK: Russell, any quick reactions sitting on the West Coast, the Bay area?

WALD: There's a lot of hype, that's for sure. Try driving around San Francisco and see all the billboards. You know, there's a lot of promise that people talk about here on the West Coast and, that's generally, capital field. But, I do think that there will be some interest. I think we're continually surprised by the technology. Every year there's some kind of unique surprise that we didn't expect, and there's this acceleration point. But I take what Daron is talking about is, you can have this unique power or powerful technology. How are you able to adapt it and make it work for you? And I think that that remains a question mark.

KORINEK: Yeah, yeah. Thank you. Let me shift a little bit more towards the medium and long term now, and I think, I can hear this in this conversation that we are having right now, too. But, basically, whenever I speak to technology experts, I get this vast spectrum of opinions. And, what it tells me is, I suppose, that there is a lot of uncertainty about our technological future, about how fast I will advance, in which exact directions will advance, and where it will lead us. And, what I oftentimes advocate in the face of this uncertainty is that we need something like scenario planning. So, if you talk about the medium to long term and you listen to all the scenarios that technology experts are mapping out, some of them include, like what Daron has alluded to before and which Daniel has discussed as a serious possibility, this notion of artificial general intelligence of AI that may be able to supplant everything that human workers can do. Those scenarios look much more like the world today, just a little bit better. And of course, there are also plenty of dystopian scenarios. And, our probability distributions on all this may, perhaps differ a bit, but I think we all appreciate the vast uncertainty. So I wanted to start with Daniel this time and ask you, what would be like an aggressive scenario look like to you, where we see really rapid advances in AI, and what would they imply for, perhaps a future without work, like what you describe in your book?

Susskind Yeah. I mean, I, so I think it's, you know, the, the uncertainty you describe is exactly right.

Nevertheless, I think it's useful, as you say, to, to think through what those scenarios might require. And, you know, you don't need to imagine a world where there's no work for human beings to do full stop. You know, you only need to think through sort of economic history and look at times where, you know, unemployment rates have been 20, 25% to know how disruptive that could be to the social order. I mean, one of the things I found most useful about thinking through these scenarios, these sort of possibilities, is where it raises, where it reveals, shortcomings in our understandings. So, one example, I spoke very briefly about the relationship between work and meaning. It's often said that work is a very important source of meaning and purpose. What's interesting, though, is that, of course, for many people, it's not; for many people, work is not a source of meaning and purpose. People would rather not work at all and get an income from elsewhere if they could. That's not to say that's a necessarily desirable outcome, but it's just, you know, for many people how they would describe their relationship with work. Similarly, if you look through history, you can find societies with very different relationships between work and meaning. Why do I say all of this? Well, because it has practical implications for thinking about the sorts of interventions that we might adopt in a more dramatic scenario. So, if you if you think about the disagreements between those who support, say, a universal basic income in a world where there's not enough work for people to do, the idea that you ought to give everyone a basic income irrespective of their, you know, position in the labor market compared to a job guarantee scheme where you, you know, actively create work in the absence of, you know, a labor market producing work. What's interesting is that actually, I think the disagreements between the supporters of those two interventions actually have less to do with the economics, the kind of technical economic merits of the two schemes, and far more to do with tacit assumptions about the nature of the relationship between work and meaning. You know, those people who support a job guarantee scheme are often carrying a hidden assumption that the reason they want to create that work, the reason they want to guarantee a job, is because a job is an important source of meaning and purpose and structure and fulfillment, whereas, you know, conversely, those who support a universal basic income, you know, often the reason they're doing that is because they think that you can find meaning and purpose elsewhere outside of work. And so actually, in deciding what kind of, you know, in weighing up those sorts of more radical interventions, we need to have a better understanding of the relationship between work and meaning, it seems to me.

KORINEK: What policies would you advocate, Daniel, to kind of robustly prepare our world for the future, given all the uncertainty that we are facing?

Susskind Yeah. So, I mean, I think it depends on which particular challenge you're talking about. You know, I set out those three different challenges before, the sort of economic challenge of this distributional challenge about how we might share our income in society if the labor market is less effective; the problem of power, what we do about the growing political power of the large companies who are responsible for developing these technologies; and then these this issue of meaning that I just mentioned before. Each of those require very different, very different sorts of interventions. You know, on the first the argument of, of, my book, "A World Without Work" is that, you know, the labor market is less effective at sharing out income in society, then the state has to take on a larger, redistributive role. And I explore the sort of, you know, the tax and expenditure implications of that, how we might tax the more valuable types of capital in the economy if human capital becomes a less valuable input into production relative to other types of capital and so on. Yeah. At the other end, you've got this issue of meaning and purpose, and that raises, you know, really, I think some of the most provocative ideas that I've written about. You know, we have a whole suite of labor market policies to intervene and shape people's working lives. You know, if you take seriously this idea that, you know, you might hollow out the sense of meaning and purpose in people's lives, then perhaps we also need to think about leisure policies as well. And, you know, the state in many ways already intervenes to sort of haphazardly shape people's non-working lives. You know, the argument that I make in my work is that, you know, we perhaps think about this in a more systematic, careful way, if we think a world with less work is a serious, as you say, scenario. So I think there's a whole spread of interventions. Some, you know, far less familiar and comfortable that we'd have to think through if we were, you know, really entertaining the prospect of, a world with less work. But I think it's important just to emphasize that's not where we are at the moment. You know, the challenge at the moment, as I said at the start, is not a world with less work, full stop. You know, in my mind, the challenge for now is really a frictional one. That there is work, but for various reasons, it's it's, there are reasons to think that that people might not be able to take up that work.

KORINEK: Thank you. Daniel. Daron, may I hand this question over to you?

ACEMOGLU: I'll be brief because I know we're running out of time, but I think, yes, there is a lot of uncertainty and there's much, much more uncertainty than experts admit. This is not unusual. There's a tendency for experts to be overly confident in their points of view, which often coheres quite nicely with their ideological stance and their pocketbooks. It is no surprise that those who will benefit from more AI investments are telling us about the very rapid advances in AI, but I think scenario planning would be useful. But I would love to see that as an open debate where people really put their predictions on the table and

admit that if those predictions go wrong, they are willing to entertain very different policies. So I would love to have, you know, the tech leaders admit what their predictions are. And if the predictions turn out to be wrong in terms of the productivity gains and the capabilities of these AI models, that they are willing to allow for new companies to enter and loosen their grips and allow for regulations. And, if the distributional effects turn out to be worse than they are hoping that they are willing to allow for a wealth tax, especially targeted at the tech sector. I would love to have open scenario planning that includes commitments from these people.

KORINEK: Russell.

WALD: Sure, I'll be brief as well. I think scenario planning is actually really interesting, and it could be quite useful at this time. But if I were to think of it, instead of specific scenarios, I would think, building on a bit of what it's already been said is, thinking backwards at the end of the scenario, what is the preservation of human dignity in this space? How are we making sure that human dignity is preserved, whether that's work, privacy, enablement, any of these things that can come from it? That will help us understand the endpoint and the end goal. And working from that and working in an institute like that I'm in, it's always, how are we thinking about the human at the beginning, through the development process, and at the end? What is the role of the human? And that whether that's the user of the technology, the community that it's affected by, or the entire society, in a set of concentric circles, we are consistently thinking about that. That's important. A quick thing from a policy front that I find really fundamental, is, one, coming back to what I was talking about the very beginning, public sector AI investment and what we do here. We need to ensure that there are a more set of diverse stakeholders at the table. If we don't, then our scenarios are limited for that planning, and they start to open up a whole in a whole different way if we do that. So there are things like the national AI research Resource or through the Create AI Act that Congress could pass, which is a bipartisan bill that will have a significant impact over time if we start to open up this whole other different area for other stakeholders at the at the table.

KORINEK: Thank you, Russell. And yeah, I think this notion of putting humans at the center, which I guess your center at Stanford derives its name from, your institute, is really crucial for this question. So I want to thank all three of you and all our listeners for participating in today's event. And I want to also thank you for your thought leadership on all the challenging questions that we have discussed today, because, as we have seen over the past hour, the future of work in an age of ever advancing AI does raise really profound questions, and at the same time, immense opportunities. And, I for one, I think we should take the concerns

that AI systems may increasingly be able to do many of our human activities quite seriously. But we may want to embrace this notion of planning for the uncertain, planning for multiple scenarios. And I also think that we should not fear a future in which machines can do a lot of the drudge work, at least as long as we manage to solve the distributed problem. As long as we manage to make sure that the advances in AI do deliver shared prosperity. So I for one, I would actually dream of a future where my children and grandchildren are not forced to work as hard as earlier generations have had to, simply because machines can perform most of the necessary tasks that we humans need much more efficiently than we ever could. What a gift that would be. But yeah, getting there and getting that transition right won't be easy and will require visionary policymaking. It will require cooperation, including international cooperation, and certainly a spirit of solidarity and shared humanity. But if we proactively prepare for that starting today, then I do believe that we can make this next economic transition the best one yet, a transformation that will allow us to leave the age of everybody having to work behind and enter an age of even greater human flourishing. So I think that challenge may be one of the greatest challenge of our generation, but I hope that we can rise to it. I hope that we will be able to pull together the wisdom and care, and to seize this chance and hopefully build a future where we will all experience shared prosperity. Thank you everybody.