Technology Policy Institute

The Economic Impact of Artificial Intelligence

2024 TPI Aspen Forum

Panelists:

Timothy Bresnahan, Landau Professor in Technology and the Economy, Emeritus, Stanford University

Carol Corrado, Distinguished Principal Research Fellow in Economics, The Conference Board

Michael Rosenbaum, Founder & Executive Chairman, Arena Analytics

Jeff Prince, Professor of Business Economics, Kelly School of Business, Indiana University, Harold A. Poling Chair of Strategic Management, Chairperson of the Business Economics and Public Policy Department

Moderator:

Scott Wallsten, Senior Fellow & President, Technology Policy Institute

Monday, August 19, 2024

8:10 AM-9:05 AM

Event page: The Economic Impact of Artificial Intelligence (youtube.com)

Scott Wallsten:

This first panel is on the economic effects of artificial intelligence. Uh, and now AI, of course, seems to dominate every tech conversation these days. I'm pretty sure it's illegal to have a conference without talking about AI. Um, you know, we may be riding the peak of the hype cycle, with projections that sound almost too good to be true. A recent McKinsey report claimed that generative AI could add $4.4 trillion to the global economy.

Now, when management consultants start throwing around numbers with tease, you might think that's an indicator to start shorting whatever it is they're selling. Um, but with AI, there's undeniably something real here. It feels like we're witnessing the emergence of a true general-purpose technology. If it's possible to identify that in its early stages, like many of you, I find myself using AI tools in my workflow every day, and its impact, to me, is noticeable.

So, the questions before us today are: What are the real effects of AI now and in the future? Do we see, or will we see, measurable productivity gains? How will AI reshape our labor markets? Do we know what to measure and what to track? Um, and if its potential impact is as large as some predict, what are the limiting factors?

Um, AI is different from other, other digital goods; those usually have high fixed costs and very low marginal costs. AI still has, you know, real marginal costs for asking it to do things, and that's different. What are the bottlenecks that it'll face? Is it data? Is it ener- energy? So, we're going to talk about all of those things and answer all of them definitively, I'm sure. Uh, so let's start with just what are the current effects. Now, what do we see of the effects on the economy from AI at the moment? Um, Tim, if you want to kick us off.

Timothy Bresnahan:

Yeah, sure, I'll take a crack at that one. So, I'm a technology in-use guy, so let me talk about the current state of AI use in the U.S. economy—not the technology, but the use of it. Um, extremely deep and extremely narrow. You've got a revolution in science, so that'll give us some productivity gains in a couple of decades. Um, you've got, in the commercial domain, stunningly impressive uh applications in a couple of dozen companies. Two of the, two biggest of them are the sponsors of this conference: Meta and Alphabet. Um, but it's just a couple of dozen.

And then you've got about 6% of American companies have an AI-based app that management or the CIO has heard of—6%. Very impr-, not as impressive as, you know, adding to Target matching at Facebook or Google, but impressive. And then, in the other 94%, you got somebody's got an experiment with an LLM somewhere in the company, which they find productive. So, you've got this, which is not yet very deep—so, you know, narrow but deep. You

can kind of pick your forecast, right? The shallow part could become like the deep part, and that could be one of the um greatest productivity revolutions of all time.

Uh, but then you've got to see a way for that to happen fast, um, which we'll probably talk about this morning.

Scott Wallsten:

And when you talk, talk about somebody who's in the company using it on their own, that's a real thing, though, too, right? We just don't see it. Carol, you might have some kinds of things you think.

Carol Corrado:

About, okay, so, um, I'm a, an economist who studies productivity and technology. Um, and I, let me just say a, a couple of things. Um, I, I think that there are many companies um, that are experimenting or using uh, you know, algorithms and massive amounts of data to um, bring about efficiency, efficiencies in their operations. Um, and I think we are seeing this bubble through in the productivity statistics. Um, if, if you, you look at recent years and take account of mismeasurement, um, we aren't really doing as badly as uh, the as many people say.

Um, and we're certainly—the United States is performing much stronger than any other advanced country uh, in its use of technology at this point. Um, but when it comes to workers, um, let me just sort of add—uh, I, I think Tim gave some, some really interesting numbers. Um, one of the uh, sort of sets of figures that, that sort of has um impressed me is that, um, seven months after the release of ChatGPT—uh, in other words, July of 2023—uh, a survey was done of office workers. This is the Conference Board did a survey of off- office workers and it oversampled uh, people who work in marketing. And what they found out was that more than, well, about 60% of office workers had uh, experimented with using uh, generative IE, uh, AI in their work. About a quarter of those worked at companies that had no—this was not top-down; it was, it was bottom-up. Uh, it was the initiative of the workers.

Um, and if you worked in marketing, 90% of the workers had, were actually experiment, experimented with or using uh, generative AI in their work, and they all reported —most of them reported that it improved their productivity and some even said that it improved their creativity or ideation. Um, that's seven months after the release of the tool. Um, this will not really show through in productivity statistics; um it's a, a worker effect. Um, they should be getting higher wages if they are more productive, and that's what we would expect to see.

Uh, it's really only when firms adopt the technology in a production process or revamp it entirely uh, that you're going to see the kind of productivity uh, experiences that, that we had, say, with the uh, with the use of computers. Um, and there, it was all very slow. Things sort of bubbled up. People talked about the narrow and deep that Tim was describing in the, in the early '90s. I

mean, it was truly there, there were a number of firms that were earning abnormal returns as a result of their investments in computers, and that's happening today too in terms of AI. Um, but it took a number of years for that to diffuse uh, to the aggregate statistics, and so I think we, we still in a waiting game.

Scott Wallsten:

So, I guess a couple of things. One is, I mean, its early days, and the first place we'll see it, you think, is in wages. Um, so have we seen any of that yet, or is it still too early? And second, when you compare it to earlier um, technological revolutions, how do we think this compares? And for anybody.

Carol Corrado:

Why doesn't somebody else take a crack at that? I could try.

Timothy Bresnahan:

Oh, let me take a crack at it—

Carol Corrado:

Yeah.

Timothy Bresnahan:

Nobody's going to like the answer.

Carol Corrado:

<Laugh>

Timothy Bresnahan:

Um, so, and it's only a so-far answer.

Scott Wallsten:

Okay? Right. Okay. I guess we can never have better than that.

Timothy Bresnahan:

Um, so far. The AI revolution is, I'm, I'm going to pick up on what Carol just said, basically a continuation of the computer and communications revolution that we already have. And, I mean, in even a sharper sense than Carol just said it—you know, you look at the firms that are really successful using really, really successful use in AI. Our sponsors, already large and successful and pretty purely digital before AI. You look at—and, and there's one from ano- another thing which uh, came from Christina McCallin and her colleagues on the Census data. You look at the

firms that are moderately successful, the 6%—um, they already undergone a digital transformation before AI had come along.

So, we're seeing, you know, the problems you might not have liked and the things you might have liked, both of them from the computer and communications revolution. Well, you get more of it.

Michael Rosenbaum:

So, oh, sorry, did I interrupt? I might push on that, although frankly, I'm a little hesitant to push on it since you guys have a much more macro view of the world than I do. I, I think the reason I'm here—and Scott can correct me if I'm wrong—is uh, I run a company that is an applied application of AI. You know, it is not a hyperscaler; it's not, you know, one of the sponsors. Um, but we apply AI to labor markets, right?

And I think that what we have seen, um, just in the last, you know, again, 12 to 18 months has been a massive shift in productivity gains at a, at more of a micro level. And so, I'll just give you a couple of examples of that. Um, so essentially what we're doing is we're um, using large amounts of information to predict the likelihood that something will happen, or someone will perform a particular way in a job, um, in the wake of some decisions—decision by the, by the employee, decision by the employer, decision by the community. And, um, you know, just so you have a sense, 16% of the U.S. healthcare labor market is on our platform. So that means we kind of know what they did in their last job; we have a whole bunch of information about them.

Um, in uh, software engineering, um, before being i-, what we're doing is we're predicting the likelihood that someone will be in the top 2% of all software engineers after receiving training. Um, and we're looking at folks without looking at the resume to figure out if they're, if they're likely to be successful with that intervention. Um, and so it, for software engineering, we're talking about, you know, the sort of tens of thousands of people, not the millions of people. But, um, average income in someone's last job before being identified is $25,000 a year; average income five years later is $98,000 a year.

Um, half of folks don't have a college degree. Um, I would describe that— I mean, you know, you guys are going to be much better than I am at sort of articulating the, the macroeconomic impact of this, but, you know, that's a pretty dramatic productivity gain. Um, I'll give you another example. Um, between the beginning of 2019 and the end of 2023, a little over 800,000 people applied to jobs um on our platform in senior housing. So, for those of you who are not familiar with the senior housing industry, it's like continuous care—um, so you know, sort of like light, sort of post-acute healthcare, so not a hospital system, something else. Um, in 2019, 49% of folks did not have previous experience in senior housing or healthcare; in 2023, 63% did. So, what that suggests to me is that these technologies um, enable folks to move across industries, um, in a way that unlocks talent and, in turn, unlocks productivity and overcomes barriers in

labor markets that otherwise prevent that productivity unlocking. So, but again, you guys will be better at thinking that through.

Scott Wallsten:

So, it almost sounds like, on toward this side of the table, it's—uh, you…I'm putting words in your mouth; they might not be right—uh, that maybe this isn't a general-purpose technology or at least we can't say that yet because it's too early. And Mike, you're saying that the effects may be bigger than we're expecting. But then on, on then on that end of the table, uh, how scalable or generalizable is it? You know, are you uniquely able to find this small group of people that would otherwise be overlooked, or, you know, is this economy-wide, that, that there's just massive numbers of people who, who could be helped this way?

Michael Rosenbaum:

I mean, now I'm getting way over my skis. I mean, full disclosure, Scott and I did work together 29 years ago, so I do have a little bit of a sense of the big picture. But, um, but I think the kinds of use cases that Arena is applying exist in a lot of places in the economy. Um, the problem of information, um, and information asymmetry, to use sort of a more academic term, but, you know, but sort of more generally, the barriers that exist to unlock, you know... I, I hesitate even to use the word productivity because you guys will be able to talk about it in a more sophisticated way than me.

But to unlock opportunities that were difficult to unlock, the technologies we're talking about have been around for a long time. I mean, I've been doing this work—I mean, my field was empirical labor economics as an academic—um, you know I've been doing this work for 20 plus years. Um. So, there are two things going on. One is the nature of the technological infrastructure—talk about digital transformation: the evolution of the technological infrastructure that makes data more practically usable is, you know, there's obviously all the stuff that Mark was talking about yesterday, which is, you know, the sort of massive increase in power that allows us to process this information in a much lower cost way. Um, you know, but then there's just sort of, I, my own view—and I'm an entrepreneur, so maybe I'm too optimistic—um, but is that we hit a tipping point with ChatGPT.

Um, and so, I'll just give you an example of this for us. So, our sales cycle in 2022 um, was eight months; so eight months from when we identified a clump potential customer who wanted to buy to when they signed a contract. Um, it since the fourth quarter of last year, our average sales cycle has been 64 days. That is a dramatic change in the way in which think about, people think about these technologies, the way people are willing to engage with these technologies, and frankly, it coincides with the massive increase in the power of these technologies.

Timothy Bresnahan:

So, my macroeconomic perspective here is, that's terrific! I wish we had another 200 of you.

Michael Rosenbaum:

<Laugh>

Timothy Bresnahan:

Right? I mean, the thing, things that are a prediction problem, right, were the, were the beginning of the applications of uh AI. Now, what ad would I click on on Facebook? From a macroeconomic growth perspective, that is not the prediction problem I would have picked to do first, but it was the most privately valuable and it was also the most visible as a prediction problem. Now, people—just not enough—are finding other things which have an element of prediction in them. I think that's a, that's one really big positive thing, but that's spreading out so slowly. Uh, Carol, do you want to say something?

Carol Corrado:

No, I just—I mean, you've really hit upon sort of a key thing here. Scott asked, what's it going to take uh to see these applications at scale? And to the extent that we are seeing AI and data being used to create efficiencies in production processes, in marketing campaigns, in supply chain logistics—all these things that are very specific to a firm—um, I think, you know. many, many firms can and want to do this, but they can't easily replicate what their competitor has already done. Not because of their inability to use the technology, not because they can't, in essence, copy the technology for essentially free, but because the data that is used to generate the outputs that lead to the decisions that create the greater efficiencies needs to be amassed again and again and again, and that's costly. So, when we have like technology bubbling up through the economy, it's because people can more or less costlessly copy an innovation. And AI innovations, the specific applications that they result in, cannot easily be copied if the data that they're based on is proprietary. To the extent it is not proprietary, then you can imagine it spreading very, very fast.

Scott Wallsten:

So, let's come back to that question of the data and the cost before we move on. Jeff, you had the misfortune, maybe, of having to uh edit a special issue of a journal um on AI. Uh, so you've just now seen probably the most current parts of the thinking on these debates. Um. What, what have you seen people argue?

Jeff Prince:

No, that's great! I, I was just going to jump in about that. So, uh we just released a special issue at the Journal of Economics and Management Strategy on artificial intelligence and the business revolution, and uh, you know, some of the topics we're covering right here are quite well covered in that issue. Um, In fact, Tim has a great piece, and I think he talked a bit about some of the insights from that; I highly recommend it. Um. Also, mention this paper by McElhin, Brelon, and others that talk about kind of the rates of adoption of AI and that it really has been quite top-heavy up to this point, at least for what they've seen so far.

The other thing I was going to highlight is, the issue really kind of paints a picture very consistent with the conversation here today. Um, and by that, I mean, you know, a lot of the research on this, of course, is very recent; a lot of the developments are very recent in this space, and the research we've got some pretty good research that's showing productivity gains in kind of specific areas. So, for example, Brelon and others have a paper that shows productivity gains in customer support. Uh. we see a lot of research in healthcare sector, um, so seeing productivity gains there with, uh what is it, medical chart coding, also with diagnostics.

Um, and then at the macro level, though, there's some mixed evidence, right? So, there's some research that's shown at the macro level that there's productivity gains, but it's kind of using a coarse measure, basically looking at sales by companies and seeing how that relates to AI adoption by those companies. But then there's other studies that uh, are surveys and, and really finding conclusive evidence. So, the, the evidence right now is kind of mixed at the macro level, at least from what I can tell. At the micro level, I think we are being able to identify places where you can really see some productivity gains in some specific areas.

Scott Wallsten:

Where do you stand on the GPT question? Is it one?

Jeff Prince:

<Laughs> I tend to believe so. Um, I, I don't think it's been unlocked yet to actually manifest as one yet, but I think it has that potential. Um, you know, I think, you know I learned a lot from Tim's paper. We have another paper by Agrawal, Gans, and Goldfarb to talk about, um, you know, thinking about systemwide adoption versus kind of, kind of like Tim's talking about, more narrow adoption in companies.

Um, and it—and you know, I don't want to spoil it, but Tim's paper really kind of identifies reasons why we might be a little worried about how well this is going to proliferate moving forward because of some impending challenges. But to the extent that those get overcome, you know, it's always hard to tell, uh, sitting here, if those—you, know, what will change, what's kind of the second order of the arc. Um and so, to the extent that there is going to be some push through on some of the challenges and getting it to be adopted more systemwide at companies, there I, I see some very, very large upside. Um, but the question, I think, is, you know, how hard is that going to be? How long will that take? Um. You know, what's that going to look like if and when we get there?

Michael Rosenbaum:

I might, um, just give a frame at least how we would think about ChatGPT. Um, and I, I think this applies: ChatGPT is—although I think it did create a little bit of a tipping point—it is part of, part of sort of a broader issue of lowering barriers to access to essentially data and technology. So, you think about what ChatGPT is; I mean, ChatGPT wasn't new— ChatGPT was just the

next version of this. Um, and, you know, I mean, we had been using some of this stuff for a long time, but, you know, but we're a pretty technical team.

Um, and so the question is, can a consumer or someone who's not data or technologically sophisticated access the power of these technologies? I mean, all ChatGPT is, like, you know, predicting the likelihood that, like, the next phrase is the right phrase. Um, and while ChatGPT is an incredibly powerful thing, I don't mean to make it too flippant, but the idea of lowering barriers—which allows the proliferation we're talking about, the proliferation through the economy of access to the power of this data and these technologies—is at its core what ChatGPT is doing. And, and again, whether ChatGPT stays in the lead or someone else is in the lead, or but it's really just sort of part of a broader continuum, but also a tipping point in the access to these powerful technologies that have been around for a while, um, to a much larger portion of the economy.

Timothy Bresnahan:

Could, could I riff on that for a second? Because I think that's a really important point. The and, and there's two uses of access in there that are really important. One is that you can make data-intensive applications at the firm level more useful because more people can use them: customers, suppliers, ideally, and employees. Um, and you know, why, why was the PC a valuable technology? Exactly—it was a moderately valuable technology when it did individual productivity apps, but then when it became, in the '90s, an access to it had the graphical user interface, anybody could use one. There were a lot of them around, and there was access to systems. It was a real reliever, and we got that blip of technical growth, technical progress, in the in economic growth in the '90s.

So that's one sense of access: more people can use it. The other sense of access that I think is really important is that, um, a little bit about what Mark was talking about last night. Uh, you know, can we find a way for more companies that do not necessarily have 87 really smart uh, IT guys uh, to use these technologies? So, that's the whole cloud thing, right? That's the embedding of uh, AI—especially generative AI—into every piece of enterprise software you might want to buy, right? That's the other, the other thing, which is much more universal than you have an app which is really good is somebody's using Salesforce somewhere, and boy, those people are working day and night to make it easier. We, we have um—you can see it in the, in the wage figures—we have a gigantic shortage of super smart people. I don't mean IT people; I mean people out there in ordinary companies. Uh, and if we can make it—if it could be like the GUI—we could just get more people and more companies using it, hey, that would solve the narrow but deep problem in a good way.

Carol Corrado:

Yeah, um, I think what, what Tim is saying is really important. If you think about what AI is and the different types of AI and the progression and the sophistication of AI as a technology, you,

you would almost say that it will have failed if it is not easy to use by a—mo, you know, an educated workforce. Um, I mean, it, it—I mean, I don't—I don't end of sentence. Um, it's just if it, it emerged in the '80s and everybody was really excited about it, and then it fell off because nobody knew how to use it. Um, and I think things are very different now.

Michael Rosenbaum:

I think one of the most important words used in that sentence is "educated." The question is, what is an educated workforce? Um, and you know, I think the classic way we think about an educated workforce is, you know, someone who knows something about this stuff, you know, has a college degree, has a—But if you think about the percentage of the United States that doesn't have a college degree, um, you know, which is an extremely large percentage of this country, as I'm sure everyone knows, um, and the implications to the acceleration—I mean, I'm saying something you guys are in policymaking, so I'm not saying anything you guys don't know—but the acceleration of the sort of changes to the economy through technology over the last several decades and what that means for basically the Sharian sort of creation and destruction of jobs and industries, and how that has now moved to less than a lifetime, um, and the political and sort of core societal stability, stability implications of that.

And, you know, and so how do we think about, again, what does it mean to be an educated workforce? Um, so me, I'll give you a really micro example just in sort of our world. Um, so we generate a whole bunch of predictions about what's going to happen. You know, the reality is that that required, until a few years ago, that required a sophistication about data. You know, it would be hard to, it would be hard to really interpret some of what we do and really get full use of it if you didn't have a pretty good like handle on math. Um, the interface of a large language model now makes that usable, um, and interrogable by, you know, someone who is working as a phlebotomist in a hospital, who is not the profile of a person that we normally think of as the sort of educated workforce when it comes to AI. And the implications of that um, and again, from, from a policymaking perspective, thinking about the implications of that and what it means to essentially facilitate and support the transition of, of the economy and workforces when the question of what it means to be educated—and if we define it in such a narrow way that it only applies to 10% of the United States, then we end up with a much more serious problem on our hands.

Scott Wallsten:

So, we have a lot of questions on Slido, um, and lots of them on labor markets. So, let's go to some of them. I'll try to kind of group some of them together. Uh, I think a collection of the questions is, um, aren't there a lot of negative aspects of this too? And there you can see the more specific things there on Slido. So, AI could reduce wages or increase them, and so Carol said that we would expect to see some of this reflected in wages as people use AI more, and so that's good for the people who are using AI. But what about the ones who aren't? Is you know, do the wages go the other way for them? I added that part on. Um, and, uh, let's see, what were some of the

others? I guess, you know, the question is, what are the negative effects on labor markets? Um, does it benefit different groups disproportionately, benefit and hurt different groups disproportionately?

Timothy Bresnahan:

I—

Michael Rosenbaum:

Go ahead.

Timothy Bresnahan:

So, I only kind of see positive effects on the labor market. Uh, the uh, you know, think, think, think about the class of people—so, you know, this is like a lot of Americans who are a screw up in high school but whom you could ask to take responsibility for the output of their job rather than just, like, literally doing their narrow job. That's a lot of people, and boy, they are undercompensated and underutilized today. If, as an access technology, generative uh AI could help us get any significant amount of value out of those folks, A, that's, you know, that's great for income distribution, and that's great for aggregate economic growth.

Michael Rosenbaum:

I mean, you know, you have a point. Um, I think, think that in general it benefits the folks we want to benefit, which is to say, you know, again, I think sort of to Tim's point, folks at the lower end of the income scale, um, I think end up being benefited by this. Um, you know, if you think about, um, you know, again, the person who's working at a Taco Bell who would be a great nurse and the sort of ability to use this information to move across the— I mean, not this super self-interested use case, but like, but that's a that benefits folks at the lower end of the scale. You know, there are, there are—you know, we're not reducing the number of nursing jobs, um, but we may reduce the number of doctors.

Timothy Bresnahan:

Yeah, yeah.

Michael Rosenbaum:

Um, because the reality is that I would far— I mean, you know, a piece of AI is probably, probably going to be better at diagnostic work than a human doctor. Um, and so the folks who end up being impacted the most are elites. Um, sorry, go ahead. You're going to be better at this than me.

Carol Corrado:

No, no, I mean, you anticipated what I was going to say, which is that we normally think of automation as, um, rightly or wrongly, as, um, eliminating jobs and eliminating low-wage jobs. Um, by contrast, particularly generative AI or types of AI will eliminate, um, high-wage jobs. Um, and he just gave a few examples. Um, so—and because AI has elements of automation, it has elements of, of, you know, diagnostic uh ability. Um, I think, you know, it's, it's going to be—it depends, but I think this is one of the few waves of technology where, um—and I may be wrong, uh, because my, my economic history isn't that deep—um, where, uh, you, you will see, um, certain segments of, um, I wouldn't say middle management jobs, but even upper management jobs and highly skilled jobs like a diagnos-, diagnostician, um, you know, not be needed so much.

Scott Wallsten:

So, is—um, some of the—so what, what explains some of the opposition to using AI? A lot of its kind of—the, the arguments are that it's going to hurt people on the lower, um, end of the socio-economic spectrum, but what you're saying is that, you know, people who are threatened are more likely to be higher up. Is this, we could—

Carol Corrado:

I think—it's a new thing.

Scott Wallsten:

So, but so is it—so are people misdirecting the argument because they're worried about their jobs?

Carol Corrado:

Well…

Scott Wallsten:

I mean, this is an impossible-to-answer politically question.

Jeff Prince:

Yeah, well, first let me, uh, just to create anonymity up here, I can say I also agree that, uh, you know, I, I think this—maybe to answer your question, Scott, I think part of it is there's this historical look that, you know, we think that, you know, that's where it's often been the case that we see the impacts, and so that's kind of the natural starting point. But I think as you think it through, it becomes clear that this could very well be different.

Um, and you know, we a lot of times people focus on healthcare for obvious reasons. Uh, I can say I know a lot of nervous accountants too, right? Uh, you know, and for good reason.

Carol Corrado:

And finance.

Jeff Prince:

Yeah, and finance, right? Um, and for good reason. And I guess the frame I always like to put on it—um, so I teach econometrics; I've been doing that for a long time—and I always try to distinguish what I call passive prediction from active prediction. And this is what I'd call a passive prediction technology, which means basically that you, you take the information as it comes in, and then you make the best prediction you can based on what you might call pattern recognition. And you contrast that with other things where active prediction would be things where you're finding structure in the world, kind of understanding what the rules are that govern what you actually observe.

And, you know, a lot of jobs that are on the lower end require exactly those kinds of skills. You might even call some of those skills common sense skills or just being, you know, able to personally engage with people. Um, and so, yeah, I, I think there is something a little bit different here, um, that kind of changes the conversation because of the nature of the kinds of tasks that this technology can take on.

Michael Rosenbaum:

I think, um, the other risk point—uh, you know, again, not to be the entrepreneur talking about the negative—um, is I think the world is just beginning to get up to speed on these technologies. And the reality is, if you do them wrong, you can create—you can cause damage. Um, and so, you know, the most obvious one is race and gender bias, and you know, and the reality is that what— you know, the way these models get trained is off of things that have happened. And so, if there are issues in situation X, the question is, how do you apply the technologies correctly so you don't end up institutionalizing bad things that have happened in the past?

And that's really where, um, again, everyone in this room, you know, really has stuff to do, right? I mean, there's—I think that is something which, um, is difficult for a private entity to do on it—to sort of control the market on its own, right? And so, the question is sort of what is the regulatory infrastructure you build that, you know, that essentially doesn't try to sort of second guess the technology? Because we all know that regulators will be behind but manages the guardrails of that risk.

And the risk—the most obvious one is race and gender, but the other risk is, you know, issues related to social capital. You know that college degrees—we assume a college degree means X; does it actually mean X? Does our assumption that a college degree means X, institutionalize certain behaviors in the market that actually aren't efficient? You know, and so how do you, how do you build an infrastructure that, that sort of manages that risk? And in managing the risk, actually, frankly, helps accelerate the development of these technologies in a positive way.

Scott Wallsten:

Scott Wallsten:

Right. So, I mean, the question is whether these new, new ways of AI-powered ways of choosing employees is better than the old way of choosing employees, which was, you know, having racist managers decide who they—who their friends are, right? Um, and so, but then the advantage of the AI is that if it does yield these, you know, bad outcomes, you should be able to fix it, um, because it's observable. And so, Mike, you've been doing this for a long time. Do you actually see changes in the way algorithms have, have been able to kind of work through these issues? Like, you know, to stop hiring every white guy named Steve?

Michael Rosenbaum:

Totally, and actually, I mean, there's sort of academic work on the way you do this. Um, there have been fewer examples of it applied, um, in actual use cases. Um, you know, frankly, we happen to be one, but we're not the only ones. Um, but, you know, but so we'll measure basically the demographic and other impacts of predictions we're making. Um, and then we, you know, we essentially are managing that real time and are tracking, you know, what are the decisions that happen in the wake of that and how does that show up in demographic decisions? Again, that could be something like race or gender, or it could be veteran status, or it could be college degree, or it could be any number of things. But, but we're tracking that every minute.

Um, the regulatory infrastructure on it is immature. Um, but Scott, to your question, um, our ability to have an impact on that—we've been, we've been able to have an impact on that for a long time. I've been doing this work a long time; we've been able to have an impact on it for a long time. Um, our ability to impact that, I would say, has had a step change in the last four years. Um, like we can—we can reduce bias multiples of what we could. We could, we could always reduce bias; we can reduce bias multiples of the magnitude that we could do it five years ago.

Timothy Bresnahan:

So, could I just say, as a comment on the public debate, that that's a discussion of a real problem, right? And that discussion of whether these technologies are either infinitely good or infinitely bad is perfectly useless.

Everyone:

<Laugh>

Scott Wallsten:

Uh, do you want to elaborate?

Timothy Bresnahan:

So, okay, so I just characterized 95% of what you could read in the newspapers or the magazines uh, over the last 12 years since AI took off. It's you got this argument among the CS guys and the statisticians: either 51% of them say it's infinitely good and 49% of them say it's infinitely bad, or the reverse. Useless for any human purpose. Then you get real problems of actual downsides that you might want to work about, which are, by the way, a little more detailed than "infinitely bad." You can work on those. Yeah, you know, it's so easy to overreact to the public discussion because there's nothing in it, right? No offense to all my friends who do that for a living.

Carol Corrado:

No, I mean, an example of that is everybody was talking about how the hallucinations that were, are very frequent in ChatGPT just render it useless. Um, or isn't that just all very silly? Well, I mean, we've just heard that one of the um, most successful changes in business processes that generative AI has had has been in customer uh, after-sales customer care. Um, and what pe-, what people have done—and I think there's a firm that specializes in this—is they, you know, take all the manuals for some product, um, they feed that into a large language model, and then when customers have questions, the answers come out of this large corpus uh that, that no one single person could ever have command of. And when there is a question that, that they do not find an answer to, they have trained the model to say, "I don't know. Let me connect you with an expert." Um, I mean, it is possible to set the parameters of these things very tight so that, um, it cuts off um bad outcomes.

Timothy Bresnahan:

The…the—

Carol Corrado:

Um, and, and, and somebody's just made a lot of money with this very simple idea by selling this software to all kinds of companies to manage their, their customer relationships. I forget the name of the company.

Timothy Bresnahan:

The, the downside of all modern post-2010 AI is it's statistical.

Carol Corrado:

Yes.

Timothy Bresnahan:

It's going to have false negatives and false positives. This limited the original, really high-value applications to things where the stakes are small, like what products does— what four products does Amazon show you? Uh, right?

Carol Corrado:

Because if they make a mistake—

Timothy Bresnahan:

They can make a mistake; your choice comes in and cleans up the mistake. Now, it turns out, other than in targeted advertising and in uh, uh low-price retail, it takes some inventive work to figure out what is the structure which deals with the risk problem by lowering the stakes. That's real invention. There's a lot of it going on. Now, is it going to happen, you know, in the '20s? Uh, that's a hard, because that's, you know, there's just not enough of it. But it's application invention that deals with both the downsides and the upsides of the techn— sur, surely, really, the statistical thing and the ability to use basically otherwise useless data is the big upside.

Scott Wallsten:

So, Harold, Harold Feld asked the question that specifically about—uh, his example was, um, tax and, and, and pharmacists. But the, his big question was, you know, what about judgment? I mean, you know, we know Avi Goldfarb work and Ajay Agrawal, but you know, AI is prediction, and judgment is for humans. Uh, where do we stand on, on, on that now? Is, are we trying to have AIs move, more into the judgment phase, and is that likely to be more problematic?

Timothy Bresnahan:

It's an organizational problem. You got this right: an organizational problem; solve slowly. Uh, I mean, as the example Carol just talked about.

Carol Corrado:

Yeah, well, another one is, uh, you know, oh there's factories that, you know, run on a lot of feedback —analysis of feedback data to sort of tweak um the machines to be very specific. Um, but, um, it used to— I mean, those, those systems were initially set up to actually then tell some human this dial may need to be changed, and then the human would make the decision. And, the goal was to run these systems and run these systems, test what the human did as, as compared with what the prediction was, and then the companies would decide, "Well, when the algorithm tells us to switch this dial, we're going to let it do it automatically, because 99.9% of the time we have found that that is the right thing to do." Um, so it, it—how long did that take? Um, I mean, this is an example from Siemens, actually. Uh, years. Um…

Timothy Bresnahan:

Decades.

Carol Corrado:

Um, well, but I meant in this one particular dial. <Laughs>

Timothy Bresnahan:

In the, in the aggregate, decades.

Carol Corrado:

Decades. Yes, exactly. Um, so, um…

Scott Wallsten:

Well, so let's—this top-rated question here. I mean, we have a panel of economists, so probably can imagine the answer, but, um, what's your take on efforts to preempt AI by protecting some jobs over others? For example, proposals to tax automation.

Everyone:

<Laughs>

Michael Rosenbaum:

People want to know.

Jeff Prince:

It doesn't seem founded at this point, right? I mean, with what we're discussing, uh, what do we expect to be the labor impacts of this? That doesn't seem to be a well-founded approach at this point.

Timothy Bresnahan:

I, I got a lot of stuff I'd like the IRS to be doing, which is not writing the 15,000-page definition of automation.

Carol Corrado:

Oh yeah.

Michael Rosenbaum:

Yeah. I think it actually ties to the prior question, which is the word—the question of judgment. Yeah, and you know, when we talk about the judgment of a human being, I think that, in general, we overestimate our ability to be good judges. Um, and I think this shows up in, you know, all kinds of ways that are malicious. Um, you know, like frankly, I, I'll use the example of when Scott and I work together. I got that job because the guy who hired me had gone to the same graduate school I did. That, that was why I got that job. Like, and I would say that, you know, there's maybe some data that backs that up, but not a lot. I mean, I think that's probably a pretty inefficient way to make that decision.

Um, and so the question is, you know, where do we believe human beings' judgment is effective and where do we think it's less effective? So, I'll give you an example of effective from healthcare. So, the reality is that a human being can connect emotionally with another human being in a way that technology cannot—certainly not today, and I can't imagine anytime, anywhere in the foreseeable future, ChatGPT mimics it. But the reality is, like, you, I have a, I have a connection to another human being that I do not have with a machine, and because of that connection, there is trust. And because of that trust, there is information that flows between two people, and that's information that's going to be very difficult for a machine to collect.

And therefore, in that situation, you know, a nurse who is sitting by the bedside of someone in the last minutes of their life is going to have information, um, that a computer will never have. And so, to try to, to replace that with statistics or technology is a mistake. But, you know, but there are a whole range of other decisions, like the Siemens example, you know, where the reality is that the human judgment isn't able to process the magnitude of information that allows it to make as effective a decision. Which is why, you know, my own view is that doctors specifically, you know, over time will play less of a diagnostic role because a doctor is going to have a bias in their diagnostic process, whereas, you know, whereas these technologies are going to be able to process more information more effectively.

Scott Wallsten:

Um, so we have about 10—well, not quite 10 minutes left. Uh, let's go back to the conversation that Carol started with, with, um, the data. And so, you know, what are, what are both the obstacles to AI development? I'm trying to incorporate a whole bunch of questions here. Um, are some of these things, like energy, are they, are they negative externalities of AI that we should be thinking about because they have environmental effects? Um, and are they not being included in anybody's production function? Um, yet, uh, it does the, do too few companies have access to too much of the data in a proprietary sense? Does open-source, uh, do open-source models help remedy that? So, okay, so all of that. <Laughter> Go. Carol, you started with the data question, so why don't you—

Carol Corrado:

Um, well, let me just give you some examples of how policy could, um, promote uh, data sharing in a way that will, um, lower the cost of, uh—or, sorry, en- enhance the diffusion of AI through the economy. And, and this is easy because you can look at some past examples that have to do with computerization and data analysis, um, where it would have only been more costly in the past than it is now. Um, so one of them is, um, credit scoring. Um, this was an industry uh, initiated uh, response to the need for better financial information on people who apply for loans, and, um, that information is shared across the industry. Carfax—um, that was an industry initiative to gather public data on accidents and the cost of accidents and attach them to the cars, the VIN number of the car that had the accident. And this greatly improved the information in the used car market.

Um, so these are examples of industry data sharing that enable quantitative inputs into, uh, in this case, risk reduction. Um, the other area we don't have so much in this country is what they call open banking. Um, this is available in many, many countries where it's actually mandated that, uh, financial institutions, uh, share individual, uh, customer information that makes it easier for one financial institution to compete against another, uh, and enhance their services. And it's actually been shown in research that there's more fintech and more investment in fintech in countries, uh, that have open banking, uh, stipulations.

So, it just prevents the lock-in effects. Um, so, you know, you're not going to get public policy to address, I think, marketing information because it's not regarded as a, a, a market failure that the government should, um, correct. Um, but, uh, the more the data that can be the input into new business models is either provided by the public, uh, through increased data collection, um, or, uh, shared um, voluntarily, um, where maybe there are tax credits for that or something like that. The more you're going to get the information being disseminated that new business models or the competing business models can be set up. I know that sounds really abstract, but, you know, we have Lyft. I mean, Uber pioneered a, uh, a business model based upon actually public data on transportation. Lyft was able to copy it because they had access to the same data. They had to develop their own software, uh, but they did that. Um, so, but you can't copy somebody's supply chain.

Scott Wallsten:

Okay, but, yeah, okay, Laura gets one question, and then we've got to wrap up. Um, did we have a mic, or you could just talk really loud? Just talk. Yeah.

Laura: Um, so the topic is economic impacts of Gen AI—one of the things a lot of my companies are using Gen AI to write the first draft of code, and they're saying we're going to accelerate our product introductions by two or three apps because our engineers don't have to edit code that these big LLMs create. If we end up with new winners and losers, but it doesn't add—does that have public policy implications, or do we not care because we're capitalists, and so the fast guy may win, and we don't care that it's a different guy than the slower guys who might have been a little more... and I'm not talking about hyperscalers; I'm talking about smaller companies that are—because we are funding—Wall Street is funding 10,000 startups using Gen AI; they're not always the first.

Timothy Bresnahan:

So, like everything else in IT, Gen AI makes, uh, coders more productive. Going back to 1952, uh,

Carol Corrado:

When the computer was invented.

Timothy Bresnahan:

—no, no, when, when, uh, assembly language was invented.

Carol Corrado:

Okay.

Timothy Bresnahan:

Uh, the, uh, so that's got to be—it's got to be a good thing to reduce the costs of, of science and technology. It's just got to be a good thing. Now, does that mean that we're going to have A+ market selection of all the different firms, especially new firms that take advantage of, uh, that new opportunity? You know, I don't think we're going to have A+ market selection, and it's not—I mean, I was an antitrust enforcer for a while. I tried to improve the quality of market selection. How'd that work out? They uh, uh they improved it, but, you know, it's not a technology regulation problem. It's, uh, you know, the technology's got to be out there broadly, going to lots of people that might want to make something happen, and then we deal with the market problems as they arise, if they arise.

Scott Wallsten:

Okay, so we, we've got to wrap up now. Um, we have—there are lots of questions that we didn't get to. I'm sorry about that, but they do get, um, uh we keep them, and we can give them to Shane, who will be doing the AI working group tonight, if we may impose on you. Um, so people who want to can actually continue that conversation. And anyway, we'll have ChatGPT answer all the questions too and send that to you. Um, so, uh, thank you very much to the panel. I thought that was really interesting. <Applause> Um, and we next up have our panel on, um, digital regulation, I believe. So, uh, next panel, please come on up.