

THE BROOKINGS INSTITUTION

FALK AUDITORIUM

BPEA Fall 2026 Conference – Day 2

Friday, September 25, 2026

This is an automated transcript that has been minimally reviewed. Please check against the recording for accuracy. If you find any significant errors of substance, please let us know at events@brookings.edu.

MODERATOR

JANICE C. EBERLY

Nonresident Senior Fellow, Economic Studies, Brookings
Co-Editor, Brookings Papers on Economic Activity

PAPER: AI AND POLICY

AUTHOR: MARTIN BERAJA

Associate Professor of Economic Analysis & Policy, University of California, Berkeley

AUTHOR: LUIS GARICANO

Professor of Public Policy, London School of Economics

AUTHOR: STIJN VAN NIEUWERBURGH

Earl Kazis and Benjamin Schore Professor of Real Estate
Columbia Business School

AUTHOR: NOAM YUCHTMAN

Hayek Professor of Political Economy, London School of Economics

* * * * *

EBERLY: Okay. Good morning, everyone. Welcome back. Looking forward to a good day today. We're kicking off with a panel of three papers on AI and policy. So we'll hear from each of the three speakers and then open it up for discussion afterwards. Okay. So we'll start with Martin Beraja on why is AI so contentious?

BERAJA: Hi, everyone. I'm Martin. Thanks Janice and Jon for inviting us here. It's a pleasure. You know, when they asked us to be on the panel, we were like, "Hey, what should we talk about? What can we say that really hasn't been said before on, on this topic?" And so we started, you know, thinking about, you know, a question that for us to, you know, sort of launch a number of papers, which is why is it that we're seeing so much conflict around AI? I don't know if we have a full answer to this, but I think we'll have the start of an answer and that's what I'm gonna share with you today.

So you may think, you know, the reason why we have a session on AI today, I think is basically the idea that, you know, we think this is a technology that has enormous economic potential and indeed, you know, even in my own past work, you know, some calculations put it at maybe doubling GDP in the US or even more, you know, extravagant scenarios in our papers. And so you would think that then all should be fine.

People should be happy about this. Yet, when you ask folks around what they think about AI in the US and in the UK and in many other places, they seem to be very concerned about you know, where this technology is gonna is gonna take us. And around this concern, we've seen you know, spurring lawsuits, regulation, you know, resistance from workers, even an encyclical from the Pope.

And so that's just to say, so why is it that now we're seeing so much conflict around this technology and is there... Is it different than in the past? And so maybe, you know, what a good place to start is, is OpenAI's own statement about what people seem to be concerned about or why is it that they think are the concerns. And the quoted thing is revealing both for what it says, but also for what it doesn't say. And so what it says, you know, when it talks about economic concerns related to job displacement concentration of wealth and the like, and also about existential concerns related to the misalignment of this technology and the potential misuse for by bad actors.

And there's a sense in which these two types of concerns are familiar. They're not really new. They're not particular to this technology even. You know, economic concerns have been there since the, you know, Industrial Revolution. Existential concerns showed up also when we with the discovery of nuclear technology. And and so what the quote omits, though, is something about how this technology was built. And at the origin of this technology, there was really a breakthrough that was, in my mind, conceptual about... an insight about, you know, how we could use people's creations, their words, text, images and the like as data for training artificial intelligence algorithms.

And initially, this just meant, you know, taking, the developers really took data online that was in Wikipedia, Reddit and scraped the whole web. But increasingly this has meant taking also people's expertise as captured in the records of work in different organizations. And and what we're gonna try to argue this is that, you know, the, this taking, this appropriation of has had an economic logic and we've, you know, I'll briefly go over that, that is based on some previous work that we've done with Noam to explain why is it that we've seen this, this sort of taking.

And and then I think what is new for today is, you know, really why we think this type of taking has raised a third type of concern that really doesn't have a clear counterpart in earlier technologies, which is the moral concern. And and I think that's the beginning to start to explain why we're seeing so much conflict around AI.

And not only that, also why the conflict seems to be so visceral. Like the language of the conflict is really one of theft and violation. And and so I'll go over that today. And so let's start now. What we see in this process is sort of, a breakthrough innovation, transformative innovation that has been intertwined with appropriation in this case of data. But it's in a sense, just the latest instance of a long history of innovations in human history that has pushed us forward, that by expanding the frontier of resources that, that are available for economic use. And that's sort of the, the point of of a previous paper that that we have.

So if you go back to antiquity, you can think about whole eras of human history are defined by the materials that were brought into use: stone, bronze, iron. Later, during the industrial age, you know, the advances in mining and drilling really brought coal and oil as new re-resources into the frontier. When we moved into the 20th century, the frontier moved into the material domains. So for instance, radio and TV broadcasting made the electromagnetic spectrum a vehicle for communications and transmission. And during the digital age, so in the last couple of decades, really consumer data has become a new type of resource that has fueled the digital economy. And now with generative AI, we're seeing a new type of data, which is not about just consumer data, but really the creative works as capturing text and words and images and together with data on-- that captures some some sense of human expertise and what workers know how to do.

And so when you look at our-- this way, this is not like something completely new. We have a long tradition of innovations intertwined with appropriation. So why is that? And and now the point that we make in this previous paper is really when you think about resources that were outside the frontier, that were not being used, a distinctive feature of them is typically that the property rights over these resources, who owns them, how c- they can be used, are often very poorly defined.

So these resources have been overlooked or, and in some cases at times they were known but inaccessible. And before these innovations that revealed their value and brought them i- into use and/or within reach. And and as a result, you can think that ex ante is either not worthwhile or even impossible to really establish ownership and think about all the possible uses of these these resources.

And hence, there's really no market or no established price where these resources trade the way that those exist for resources that are already in use for which property rights are well defined. And and as a consequence is-- as a consequence what we see in very, you know, most important historical cases is that innovators just that discovered these resources went in and appropriated them. And and in anticipation of extraordinary rents from being able to keep these resources exposed and their taking being legitimized. And and as a result of this taking, there's been numerous conflicts with counterparties that had competing claims.

So let me just name a few of those. So when you go back to the history of how, you know, crude oil happened in the US, prospectors... the conflict there really, you know, prospectors went and started draining reservoir, reservoirs that spanned different property lines. So the conflict there that was subsequently resolved through a number of landmark cases was

about whether this taking was all right or not and that led to, you know, the establishment of property rights and what is now called the rule of capture in the US.

In the case of radio, early broadcasters just took frequencies from the spectrum for themselves, started transmitting, and that led to interference and eventually licensing after the Radio Act of 1927. In the digital economy, a lot of the platforms just when they discovered that they could target ads with consumer data and make behavioral predictions, they just again went in and collected, aggregated, and monetized all this user data without any even asking for permission. And eventually or even still is that led to different forms of privacy regulation and disputes over data ownership.

And now with generative AI, we're seeing again a taking in this case of online content and further expertise and a number of lawsuits by publishers and authors under to this, you know, try to debate whether this falls under fair use or not.

Okay. So we see that as sort of like the, you know, the, the economic logic for why we've seen this type of appropriation in, in, in history being intertwined with innovation when the innovation is really about bringing resources that don't have well-defined property rights. But that doesn't explain why the conflict this time is moral or why it looks so visceral.

In the past, really these conflicts were about, you know, revenue streams, sometimes about externalities, sometimes about, you know, legal rights. But when you look at the conflict over AI, it has this extra dimension that is about moral concerns over human dignity and a language that really is about theft and violation, not just a taking. And tellingly, if you look at, you know, the, the Pope's encyclical in, in 2026 the main point of that encyclical is really saying that technological progress continues alongside covert or overt violations of human dignity, and that with AI, there's this danger of translating the mystery of the person into data and performance.

So that is a very -- If you just think about what that is saying, it's a very strong language that really doesn't have a parallel in previous encyclicals. There's been two more related to technologies, one during the industrial revolution, *Rerum Novarum*. That one is basically about distributional issues and economic issues of, you know, the, the fortunes that capitalists are making and the consequences for workers. There's been another one in the mid-twentieth century that was basically about nuclear weapons and and there it's about the, you know, the, the fear of extinction and survival. So really this, you know, moral dimension - even for the Pope, really didn't appear in this, in these previous encyclicals in the way that it's appearing now.

And so why is that? Why is the language this time so visceral? We think it has to do with the nature of what was taken, the nature of the resource this time, not just that the resource was taken. So when you think about, for instance, like land, oil or spectrum, these are big things, and things in a sense they could always be taken from us, could always be sold, bought stolen. But creativity, expertise, the things that people make and know, which are now being captured in data and being used to train LLMs, to many this feels like who we or they are. And and if you think about how the LLMs were trained, this is not just copying.

It's not just like there was a-- you pirate a book and then you, you know, produce that book and replicate it elsewhere. Because when you pirate a book, the book's author is still known. It's just that maybe they're not collecting royalties. But for LLMs, these just absorb millions of

work, no trace of who made what, and learn to do what humans could do. And and so really what was taken, it was not just the royalties, it was not just the the economic aspect of these resources, but really the authorship of what humans made and to many what they are. And they believe that with that, to write, to draw, to create is sort of, you know, were things that only we humans could do. And within that, we shouldn't notice that this is not just, you know, what people mean by a violation of human dignity is not a legal thing. This has nothing to do whether or whether, you know, copyright laws were broken or not. It's something entirely different and it's a moral aspect of this conflict.

And and so, you know, in the paper, and maybe not to get too religious here, but that we call this the sort of the original sin of this technology, where-- And I think we chose the term with care because there is a taking of this kind followed by a breakthrough of this magnitude. It sort of has a meta- the parallel in Genesis, the first regression that came from the taking of the fruit of knowledge. And in this case, it was also perceived as a profound violation by God in that case. And, but yet it made everything that followed possible, like civilization and all, all knowledge. And maybe so it may be with AI that you know, promises to bear a great more deal of fruit, but it was humans this time really who were aggrieved.

And and so where does this leave us? And we thought, look, we should talk about policy at some point during this talk, not just pontificate on, on the moral aspects of this. And and I think there's a very clear policy implication here, what can work and what won't work to resolve this conflict. And so in my own work, I've actually thought about, you know, things like universal basic income, slowing down automation, retraining. And these are policies that are often discussed and, but they're designed to address the economic or the existential concerns, not-- and regardless of what you think about them, I have my own views that we can talk about them offline-- they are not designed to address this extra moral concern. And so what is, what kind of settlement should we have here? And it's not just a, a form of compensation.

And so maybe a good place to start here is what, you know, people like Sam Altman is saying. So he, you know, he... This is a quote from him where he says, "Look, I no longer believe..." He was a big proponent of universal basic income early on. They actually designed a whole a whole experiment on UBI. He's like, "I don't believe in this anymore. I much more interested now in thinking about a kind of collective ownership of the technology." And and we think he's right when he's saying that.

In that if you think about what an ownership stake means, it's really recognizing humans' role in creating the technology in a way that just a compensation, a monetary compensation or transfer does not. That address a different type of of concern. And and maybe an analogy for many of us here that kind of works is that if you think about your teachers, no, your teacher's influencing ways that you can't really point to. You're not citing your teachers, your advisors in every paper that you have. Yet your teachers, or we, when we have students, we do share in their success. How? Through the recognition that they later get that we were the one that that, that trained them. And so an ownership stake in the technology that we, where we are as teachers of this technology really restores the same relationship that we have with our own creative work, and that is now just being, you know, untraceable and embodied in the weights of LLMs.

And and so an ownership stake really is a form of what in the literature sometimes is called restorative justice in a way that, you know, other forms of compensation are not. That is all I have. Thank you very much.

EBERLY: The second speaker is Luis Garicano.

GARICANO: How it's working? I had never had a Genesis in an economics paper, so that was cool. And I realize I'm not quoting, I'm not quoting my teachers here, so I should: Sherwin Rosen. Sherwin Rosen is great with a lot of what follows.

So, I want to talk about the division of labor between firms, between individuals, and I want to make a very, very basic point. I will make it theoretically. And then I want to show you some initial preliminary evidence that seems consistent with this mechanism going. So calling the paper The Vanishing Advantage of Specialization, and I'm going to argue that the application will be law firms and professional service firms.

I want to argue that the value of specialization versus doing things in-house is gonna drop, and division of intelligence does that. So what's the value of specialization? We all understand that airline economics, when you know there's a fixed cost of a plane, the more people go into seats, the more seats you have filled, they lower the average cost.

What happens with knowledge is very similar. Acquiring a piece of knowledge is extremely expensive. Economists tend to think that, okay, if I show you my theta, oops, you know my theta instantaneously. That means everything that I know. In fact, we all know that communication is expensive and learning is expensive.

We spend hours acquiring and transmitting knowledge. And specialization is about some people acquire a little bit of knowledge and utilize it in-intensively. You acquire what you need to know, and you ask or deal with others for the rest. Now, when you reduce that fixed cost for both the specialists and the non-specialists equally, you give me a machine that requires-- that reduces the cost of cognition, the value of specialization goes down.

Why? Because the whole point was to spread that fixed cost. If that fixed cost is lower, the value is lower. I apologize for putting it in numbers. You'll have it in Greek letters in a second. So, so this is a hundred thousand dollar cost of acquiring capability. Think of a firm that occasionally has a labor dispute.

And okay, I could hire a labor lawyer, I could hire a specialist, or I could outsource it. And the firm itself is only going to use it ten times, and the specialist is going to use it a thousand times. They only do labor law, and they are going to go to all firms and use that capability repeatedly.

So, there is a cost, however, of doing a transaction outside the firm. Call it a transaction cost, call it translation cost. I need to go to a lawyer explaining this is what we have, this is what's going on in my firm, etc., which here this transaction cost is two thousand. So clearly, outsourcing is much cheaper here because I have the hundred thousand divided by a thousand, that's a hundred plus two thousand. I only spend two thousand one hundred per problem. However, I now am going to reduce the cost of this capability to ten thousand. The fixed cost is going down. Now my advantage here is ten divided by a thousand versus ten thousand divided by ten. Okay, so it's still an advantage, a thousand versus a hundred.

But now the two thousand outside cost is significant, and I don't need to incur it. So what this table tells us is even though the outside lawyer and the inside counsel are using the exact same technology, because the fixed cost of solving these type of problems is down, the

advantage of the specialist is lower. That's all I'm going to say. I'm going to show you theoretically, and I want to show you empirically what's going on.

So knowledge the firms and markets, the firm is going to be somebody who owns the local knowledge of its core activity. I know what's going on in my environment. I know my customers, et cetera. And, you know, there are many problems that I need to do as a firm. And the specialist is somebody who sells the same knowledge to many to many customers, and there's going to be a cost of crossing the boundary. So just for those who are missing the Greek letters, this is the same that I said before.

So there is-- I'm going to order all the problems like, like in much of my previous work on hierarchies. I'm going to order the problems from common to rare. And so there is going to be the number of times that a firm faces a problem per-- of a particular class is going to be Q . And I'm going to order them so that as I advance, they are rarer and rarer.

I face, yeah, Microsoft, I have lots of lawyers, but how many times am I going to do a merger? Well, since I don't do so many mergers, when I merge with Activision, I actually use an outside counsel. Even though I'm big, and even though I have thousands of lawyers, that's extr-extreme case. So there are cases that I deal with often, the cases that I don't deal often.

N is the number of cases that the client, that the specialist can pull across many pl-clients. K is going to be the fixed cost, which depends on the AI, on the machine. And τ is this cost of crossing the boundary, this transaction cost, this cost of asking somebody who does... You need to translate your inside language into other language, and you need to be context.

You can imagine all of that. And we're going to have the cost is decreasing in the increase of the technology. So like in my table, I have a unit cost of Q -- K divided by Q if I do it inside, and the unit cost of K divided by N plus the τ . My-- I don't think my thing works. Unit cost of K divided by N plus the τ .

So very simply, I keep the problem the problems that are frequent enough to have the capability inside, and I ask the specialist firm otherwise. So very straightforwardly, there is going to be a cutoff. I'm going to keep some problems, I'm going to ask the rest. And clearly, lowering the fixed cost identically, okay?

It's not just I'm telling you, "Oh, I'm going to give them a technology that replaces the lawyer." No. I'm giving the lawyer the same technology, and both of them acquire the same advantage. Lowering that fixed cost for both is going to increase the number of problems that are done in-house. Stigler wrote a paper that some of you know, that was the division of labor is limited by the extent of the market, that says when the market expansion-- when the market expands if I have a bigger market, industries, they-- it start to generate specialists, okay?

So, so IBM gets replaced by Intel and Windows and hard drive manufacturer, et cetera, because the specialists can specialize in those few. In that segment, there is enough demand outside to support a specialist, and that clearly drops out of here.

So just to put it in a graph. So what we're gonna have is we're gonna shift the boundary of what is done inside, and we're gonna have more done inside and the outside, the outsider,

the specialist is going to focus on rare-- They're going to-- Here, they're going-- and we will discuss that. Jim was pushing me for that. They're going to focus on rarer problems, more difficult problems. You can do this by difficulty the same. The specialist firm is going to have to focus on rare problems, more difficult problems, and less problems.

Now, I want to contrast this with Coase, the Coasean singularity. And the Coasean singularity idea is, look, firms are going to disappear. Transaction costs go down. You have your agent. Your agent will go around making transactions in the market. It will put things together and give it back. So the Coasean singularity says the transaction cost the trading gets cheaper, the transaction costs go down. That means that you're going to do less and hire more things in the market. You're going to do fewer things, okay? It's the opposite prediction. I'm a firm making gadgets. Now, I can transact everything with my agents. I send my agents going things around, and they do stuff, and I don't need to do anything, okay?

The knowledge story says, well, you know, you have a machine that lowers the cost for everybody, in particular for you. Why would you go and incur this transaction cost? So we actually have two effects here, a race between lower transaction cost and lower knowledge cost that can shift the boundary inside or outside, depending on whether the utilization effect that I'm arguing is stronger or the transaction cost effect is stronger.

Now, when you work through the logic of these ideas, you can see that all professional service firms, doctors, lawyers, et cetera, I will show you some evidence, get squeezed in the middle because you have kind of general knowledge and I've written a book about this that I present in Peterson in a couple of hours, which is called Messy Jobs.

The stuff that can be easily peeled out and is generalizable and verifiable gets dropped and it's done by the model. And on the other hand, the general-- the the generalist who is kind of the firm that actually makes the gadget can solve many more problems inside because that fixed cost is going down.

So if you're in the middle, you're stuck with like rarer problems the more difficult problems, and that could be a that could-- that squeezes you in the middle. So we are in very early stages, and I am not going to call this evidence, okay? I do call it evidence, preliminary evidence, okay?

I want to give you a, a few touches of things that suggest this could be going on, and you tell me what you think What's happening with AI use? What do people think about the boundary freelance markets? And now I sh- gonna show you some macro systematic data on what's happening with hirings by these professional services.

These are two OpenAI researchers, Czin and Richmond, 2026, who show that they don't have time series because there's no ChatGPT before, okay? So we don't have the time series. But they show that of occupation-specific messages that of messages that are related to an occupation, that we can classify as law, medicine, et cetera, fifty-six percent of the messages that people send to ChatGPT are outside of their own occupation.

Now, you would ask me, like, how many times does a lawyer ask about accounting in the past? I don't know. This is like, with ChatGPT, I can tell you sixty percent of the times. Okay? They ask across... Sorry, forty-three point five percent of the time. I'm reading it wrong. The

cross-occupation messages are dark. When we do occupation messages that we can classify as within an occupation, almost half are outside of my occupation.

I'm an economist, and I ask, "Hey, what happens with my law? What happens with the house I'm making? What happens, et cetera, et cetera." When you ask lawyers, the Association of Corporate Counsel made a survey, and they said, "Hey are you using legal legal LLMs?" And they say, "Yeah." Half of them said yeah. And how does it-- how does it make you feel about the use of outside providers? And they say, "Well, I don't need outside providers."

Okay, so they report, and I want to show you this is not happening yet. They expect generative AI to reduce the reliance on outside on outside providers. There is all this evidence from the freelance market, where the freelance demand for freelancers went down immediately after ChatGPT. And not only went down, but it does look like it's going towards harder tickets. So when you see which in these markets, in these freelance markets, which things people are still doing, you do see this tale of harder problems. It does look like people who are still providing this are specializing in harder and and more skill intensive.

And here's my macro evidence, the BLS occupational classifications. Just to give you one fact, between legal, accounting, software, and management they added three hundred and thirty-seven thousand jobs. Three hundred and thirty-five thousand were outside law firms, accounting firms, consulting firms, and software.

Okay? Now, this is different than before, as I will show you in a second. This is the trends. Now, obviously, it's not been going on only with ChatGPT, but I would claim that the drop in knowledge costs predates ChatGPT and is a continuous trend. So Google and databases, et cetera, all of them have this effect of lowering the fixed cost of expertise. We are all having access to expertise. What is clear is that post-ChatGPT, this trend seems to accelerate for lawyers, accountants, software developers, and consultants and market specialists. Look at graphic designers, which it doesn't happen. It's not the case that the biggest share of graphic designers are actually working outside of graphic design firms. In fact, more are working inside graphic design firms.

I want to show you what happened before. So, so basically what you see is that in the specialist firm, there used to be quite a bit of growth. This is the weird times, it's just there. These are the weird pandemic times, so things reverse, and then we have an acceleration of the trend. I'll show you. Now I want to say before I show you the actual employment numbers, I wanna say that, and thanks, Jen for pushing me to to think more of the-- this Jevons type idea. That nothing of what I said means there are less jobs for lawyers. Okay? So the fact that the share of people doing work inside law firms is smaller doesn't mean-- sorry, the share of problems that are dived at law firms is smaller doesn't mean that you need less people. These rarer, harder problems could increase employment even in the law firms. The share goes down, but the total quantity goes up. In fact, what we see is that right now billable hours are going and hourly rates are going up. These are actual employment numbers post-GPT, and these are the guys that we were surprised that these are the designers.

Ah, do I not have the designers now? Okay, so in all of them, -- this is a separate graph. Okay. In all of them, there was still quite a bit of employment growth in some of the ones that we see the professional share, the share that is employed in the firms going down.

So I have three policy things that I wanted to discuss, and it will be hard. One is licensing. These professions are the only people in the world, well, actually, well, maybe professors do that same as well, who decide who gets to practice what they're doing. And you can imagine that a huge part of the political debate over the next few years is going to be over those boundaries. I'm not that optimistic that they can really put barriers because if you are the firm and you can solve the problem yourself, you find yourself some cheap lawyer who's going to sign that, yeah, the problem is right. Even if they put the boundary, the effective supply by may, may significantly increase. But licensing is one huge issue that we're gonna see as a result of more work being brought inside and less work being done by these big professional firms.

The second is competition, of course, to the extent that the model rents are high and that the model owners are monopolist of this generalist knowledge, this K effect doesn't happen. If the K, the fixed cost is kept high by ChatGPT, and there is no competition, and we don't know that's gonna happen, then that effect of moving work inside the firm is not gonna happen. So that's the, the competition effect is one of the ones that could actually paradoxically kind of help these, these professional firms.

And third, training. Professional service firms provide a service to everybody. I have worked with Luis Rijo, which I should be citing here. So I didn't cite Rosen, and I don't cite Rijo. So I did it yesterday in the plane, and that's not the best place. Sorry. So I have worked with Luis Rijo, which is called Careers in the Age of AI that makes the following point. The point of apprenticeships is you cannot pay for the general training because they will have to lower your wage to so much that there's no possibility. The way you pay for your training is in a knowledge for work exchange. You have a dynamic contract where I work a little bit, I pay you with my work, you teach me a little bit. And the firm doses how much they teach me little by little to keep me there kind of enslaved, if you want. Sorry, but it's a relational contract. So I'm happy to be doing this. I'm learning, so I pay with my work, and they extract my surplus.

The LLM is commoditizing the currency that I use to pay for my training. It's devaluing this currency. It's saying this currency is not-- I was good at doing PPTs, but now PPTs can be done automatically, so how do I pay you to train me? Think of your postdoc, pre-docs, et cetera. So it's a problem that we discuss in separate work and how new institutions should emerge, and I'm very happy to talk about it in the break, but I'm two, two minutes away, which I'm never late, so I am so sorry.

Firms we would expect firms to do a larger share of support of work themselves, and the specialist advantage drops as the fixed cost of knowledge drops for everybody. We all become generalists. True for firms that will incorporate more of the knowledge inside and ask less outside. Thank you.

EBERLY: Third panelist is Stijn Van Nieuwerburgh.

VAN NIEUWERBURGH: Thank you so much for having me. We're gonna go from the abstract and moral from the papal encyclics to the Gaussian singularity to something much more pedestrian, data centers.

So we've been talking about AI so far, mostly sort of in very intangible terms as a, a revolution in terms of models, software, and data. But, you know, we all know that it takes a lot of hardware to produce all of these tokens. And something interesting has been

happening technologically in the sense that we now have, you know, custom-built data centers that are specifically built for these training facilities that use specialized chips, GPUs rather than CPUs. They use special cooling equipment you know, backup generators, network equipment and all inside these purpose-built data centers.

And so what's going on here to facilitate all of this compute is really a large-scale infrastructure investment cycle. And so what I wanna do in this paper is talk briefly about how large the -- what's the size of this build-out, and then also think about the financing of this build-out. And I wanna point out that there's something important that has been happening with how all of this all of this infrastructure investment has been financed over the last few years.

So to measure the scale of this build-out, I wanna start with a number, which is right now, today, we have in the United States, all my numbers are just for the United States, about fifty-seven gigawatt of compute capacity available, right? So we're not measuring this in square feet, we're measuring this in, in, in power terms. Now, if you look at the latest database of planned data center projects that are you know, either u-under construction or planned to be built, there's an additional five hundred and nine gigawatt on the table. Now, I do not think that this five hundred and nine gigawatt is a realistic number.

So in the paper I make several assumptions to get to what I think is a more plausible number of compute that might be built between 2025 and 2032. And you can look at the paper for my assumptions. I think they're quite reasonable, and I try to be conservative. And I arrive at about a hundred and eighty-three gigawatts that I think is plausibly going to be built over the next eight years. Maybe another a hundred and seventeen gigawatt that will be built after 2032, which I will not consider, and then the remainder never completed.

If you do the math, and I'll show you in a numerical example on the next slide, building this hundred and eighty-three gigawatt will cost the US economy about ten trillion dollars over those next eight years. That amounts to three point six percent of GDP. And I'll try to put that number in perspective as well.

So here's what a modern data center campus costs to build. This would be a two hundred megawatt facility. Just to put this in perspective, a few years ago, we would be building two megawatt or five megawatt facilities. So these are very large campuses of GPUs model training facilities. The real estate costs about eleven million dollars per square per megawatt. So two hundred megawatt, two point two billion dollars. The power infrastructure you know, the power to generate it, but then also the substations to bring the power to the site would be another four hundred million dollars.

And then actually, by far the highest cost is the IT that goes inside these. So that's five point six billion. That would be the GPUs and the fiber optic and all the other sort of connection equipment. So eight point two billion dollars for a two hundred-- for, you know, what's what is these days a run-of-the-mill AI training facility. So thirty percent, one-third real estate and energy infrastructure, two-thirds compute hardware. Okay? So you scale this up to a hundred and eighty gigawatt, and you get to my ten trillion dollar number.

So how large is this? Well, I went back and looked at the economic history record to compare it to previous capex booms in the United States. Looked at canals, railroads, electrification, highways, and the internet boom. And as you can see it's, you know, the it's

the largest one of them all, even relative to GDP. And it's about fifty percent larger than railroads. Okay? So it's enormous. We know that this boom is already visible in macroeconomic aggregates. It meaningfully contributed to US GDP growth in 2025 and first half of 2026. Like many of these other previous booms, AI is a general-purpose technology that has lots of downstream effects on all other sectors of the economy, creates a lot of complementarity investment, and then the returns depend on uncertain future demand, right? So there's a lot of risk involved.

Indeed, if you look at this historical record, many of these previous booms, in fact, all of them except for the highways, ended up in some sort of financial crisis, either a bond crisis, like, like with the railroads, or a large equity drawdown, like in the case of the internet. One important distinction between the previous booms and this one is that unlike canals and highways, the useful economic life of a lot of that IT that goes inside the, the data centers is relatively short, maybe four to six years.

So the second point I'm trying to make in this paper is that -- so we have this enormous investment. It needs to be financed. There has been a massive shift in how the financing of that investment has taken place over the last few years. It used to be that the five hyperscalers, Oracle, Amazon, Alphabet, Microsoft, and Meta, would do most of this financing on balance sheet with retained earnings, with cash and maybe with a bit of corporate debt. Standard, you know, senior unsecured corporate debt. What has happened in the last few years is that the amounts have grown so rapidly that they have outpaced the cash flows that these companies have made, right? So whereas you see that back in 2020, these companies made two hundred and fifty billion dollars in operating cash flow and only had about a hundred billion of investment. Today, their capex is eight hundred billion relative to only seven hundred billion dollars of operating cash flow.

So these companies are growing, but the capex needs are growing faster. And so, you know, partially out of necessity, partially out by choice, these companies have been essentially moving them, moving the financing from on-balance sheet to off-balance sheet.

So I call this the great risk shift. We now have a whole ecosystem of both equity outside third party equity investors and debt investors in this space. So we have specialized third-party landlords that built these data centers, infrastructure funds that own them, REITs, private equity funds, private credit funds as well. On the debt side, we have a lot of bank debt, both mortgage debt and syndicated bank debt. We have private credit, we have leases, joint ventures, and structured finance as well. On the IT side, you know, oftentimes the, the hyperscalers would still own that IT on-balance sheet, but increasingly even the GPUs are being financed with external debt in the form of GPU-based asset-backed securities.

So what has changed is that the users of compute are no longer the owners of these assets. That risk has been shifting towards other investors. I give an example in the paper which I think is not just an example, it's really a template for how these modern facilities are being financed. So this is Meta's Hyperion campus in Richland Parish, Louisiana. Two gigawatt, thirty billion dollar investment. Meta was initially going to build this campus itself, but it decided to change course and rather than do it on-balance sheet, it mostly financed it off-balance sheet. It brought in a private credit firm, Blue Owl, that acquired an eighty percent equity interest you know, together with Meta's twenty percent equity interest constituted the equity. And then this joint venture set up a special purpose vehicle called Beignet, which went on to set up another special purpose vehicle that issued debt. That entity issued twenty-seven point three billion dollars of debt. This was the single largest investment-grade

corporate bond ever issued in American history. It was rated A+ by Standard & Poor's, one notch below Meta's double A minus credit rating. And in the process, what has happened is that this, this thirty billion dollar investment was financed with ninety percent debt and ten percent equity.

Whereas if you look at Meta's own balance sheet, Meta has about a three percent leverage ratio if you look at market equity, twenty-five percent if you look at book equity. So we have a mass-- this shift from on-balance sheet to off-balance sheet financing is not only bringing a lot more capital into the picture, it's also leading to a large increase in leverage. So I don't need for you to look at this picture in too much detail, but suffice to say that it's complicated.

And there's this sequence of special purpose vehicles, some of which contract with the power company, some of which issue the bonds. All of this is meant to make this, this facility bankruptcy remote from Meta's balance sheet. So Meta here is basically the sole tenant of this facility. It's gonna use all of this GPU.

And in another interesting twist, Meta has signed not one twenty-year lease, but five consecutive four-year leases starting in 2029 and until 2049 when the bond, the the twenty-year bond matures. What that means is that Meta can decide to renew or not renew every four years some of the buildings on that campus. If it terminates one or more of these leases, then these assets, through a process, are gonna be offered for sale. If the sale value is insufficient, Meta is on the hook for what's called a residual value guarantee. That amount should be enough to pay back the principal on the bond. So this residual value guarantee preserves much of the long-duration credit support that a developer would require to build one of these massive facilities while offering Meta some operational and some strategic flexibility.

Now something something interesting has happened in this structure, something I call a probability vacuum. So under general accepted accounting principles Meta does not need to put these future lease commitments on its balance sheet because it's not at least seventy percent likely that they will materialize. But Meta also must not put that residual value guarantee on its balance sheet because it's also not seventy percent likely that it will materialize. One of these two liabilities will materialize with probability one, neither of which shows up on the balance sheet. Is this a big deal? Yes, it is.

There's three trillion dollars of these liabilities coming to just these four companies in the next several years, only about twenty percent of which are currently re-reflected on their balance sheet. The other eighty percent will be coming on balance sheet at some point in the future. These are these future lease obligations, as well as future purchase commitments for chips from Nvidia and Google and Broadcom and so forth. Okay? That's even without counting off-balance sheet, off-balance sheet construction guarantees with power companies that are not even reflected in here yet. Now, this structure, this complexity of setting up these special purpose vehicles, of course, comes at some cost. In fact, Beignet's debt was issued at a higher yield, about a hundred basis points about what Meta would have paid if it had financed this on balance sheet, which amounts to about five billion dollars in extra interest cost to Meta, which, you know, Meta is ultimately gonna end up paying for this in the form of higher lease payments.

So why would Meta do this? Some of it is out of necessity, right? It's sort of running out of free cash flow to finance all of these expensive investments. But some of it is out of choice. I

think these hyperscalers are trying to preserve an asset-light corporate profile so that they can be valued by the stock market as a fast-growing software company rather than a heavily loaded infrastructure firm. And they also like to preserve their pristine credit rating. Okay. So I think that's partially the motivation for this great risk shift.

So the last part of the paper thinks about the risks to this system, right? And so I think there's sort of three large categories of risk when we think about this from the investor's perspective.

First one, of course, is that the demand for AI may disappoint. Maybe it's not this incredible productivity-enhancing technology it's made out to be. But even if the compute holds up and keeps growing very fast, maybe the pricing power, the revenue-generating potential of this technology is not as large as it might be, right? I'll come back to this in the next slide.

There's also a lot of concentration. A lot of these facilities have just a single tenant, so there's what we call single tenant risk. There's a second category of risks that have to do with technological change. The AI hardware is evolving rapidly. NVIDIA is coming out with a new GPU every year. Compute efficiencies and compute modes could be changing. We might be using quantum computing, which is a hundred times more powerful ten years from now. And so these, these very long-lived investments are subject to, like, shorter-term technological change.

And third, there's a lot of execution risk, right? Power, transmission, permits. We talked about the opposition to a lot of these AI and data centers, right? So it's been very hard to get these things off the ground. And on the one hand, that could help slow this build-out and maybe preserve some pricing power. On the other hand, if you have already issued the debt and now your project is stopped in its tracks, you have a problem, right? Meta has already issued the bond. You know, if there's now a moratorium on data center development, that would be a big problem. So the timing between when the cash flows arrive and when the debts when the debt has to be paid is crucial.

So just to come back to this demand story, you know, one last calculation I do in the paper is this question of how much revenue does this build out have to generate in order to be to generate a decent return on investment. And so I'm using ten percent as a return on investment here. That's a little bit above the equity cost of capital of the hyperscalers.

It's like a reasonably a reasonably conservative number. So I'm asking the question, if we want to earn ten percent return on this massive hundred and eighty-three gigawatt build out in steady state by twenty thirty-two, what do the annual revenues from this technology have to look like, from these data centers have to look like? And the answer is three point seven trillion dollars. That takes into account that the IT will depreciate much more quickly over six years the real estate and the infrastructure the energy infrastructure over twenty years. Okay? So is three point seven trillion dollars a lot of money? Yes. It's nine point seven percent of US GDP. If you calculate it per GPU hour, it's about five-and-a-half dollars, which is roughly what the current cost of GPU is for the frontier model.

So what we're saying here is if we can quadruple the size of this data center infrastructure, and we can still charge the same price we can charge today in this heavily constrained world for the frontier model on all of that compute, then we might be fine.

But that's a lot of big ifs. For context, OpenAI and Anthropic are currently making about a hundred billion dollars in revenue per year. So to go from a hundred to three point seven trillion, they would need -- those revenues would need to grow eighty percent per year for the next eight years. No company in human history has ever grown for eight years at eighty percent per year. So the risks are more correlated than it looks. These chip designers, cloud providers, model developers, data center operators, they're all dependent on each other through a web of connections, long-term purchasing contracts, long-term leases, strategic investments. Nvidia is investing in OpenAI. Nvidia is investing in CoreWeave.

Some of this is vendor financing so that these providers can buy the chips from Nvidia leases, financing arrangements. And so there's a lot of circularity here. I'm in the process of trying to connect all of these linkages between all these different providers. This is what I spent most of my summer about. Hopefully, I'll have some results in the near future to really sort of take a granular look at all these connections between all the players in this ecosystem.

Suffice to say that it's become very hard to track who's the ultimate bearer of the risk through all this complexity and opacity. I faced this when I tried to look for all these connections. Some of them are private contracts between private parties. They're very difficult to find.

So that leads me to my last point, which is measurement. I think from a policy perspective, it's incumbent on all of us to try to measure this problem better. The financial exposure is spread around a range of entities, corporate balance sheets, leases, joint ventures, guarantees, project debt, a lot of which sits on banks' balance sheets SPVs and private credit.

It's very hard to put it all together, but I think it's clear, to me at least, that this is a source of systematic risk to the US economy, in fact, to the world economy. And so it would be good incumbent upon all of us to try to measure it as clearly as possible. Thank you very much.

EBERLY: Thanks very much. Now we have time for questions. So let me start with Arvind Krishnamurthy

KRISHNAMURTHY: Okay. Thanks. Those were three very thought-provoking presentations. Let me first ask a question to Stijn. So if I take the numbers you put up, and I'm --your probability vacuum basically suggests that one way-- the right thing to do is to take all of those debt payments in some form or the other and consolidate it back into Meta's balance sheet.

And if you do that and use your ten trillion dollar future expenditure, and you consolidate ten trillion dollars back onto the collection of hyperscalers, what does that mean for their leverage ratios? So I guess the, the background question is, should one think about the risk here as being highly lever, highly levered hyperscalers that we don't see?

Or should one think of the risk sitting fundamentally in the credit providers, private credit, et cetera? Where should we think that the big risk is coming? And then the other question that's related here is, what I didn't hear you say, which is a thing that when we think about financial fragility often pops up, is short-term debt. Everything here sounds like long-term debt, but I'm not sure. So I'd like you to speak about that as well.

Luis, I had one question for you, which is you know, your first example was decreasing cost to specialization. Does the patterns that you were documenting later allow us to figure out how much cost to specialization has gone down? Thanks.

EBERLY: So we'll take a few questions before we come back to you guys. You'll get a chance to respond. So David Byrne.

We'll get you a mic, David. I should remind everybody we mentioned this yesterday, that the conference is live-streamed and recorded, so, remember, keep it in mind.

BYRNE: Yeah. Right. So I'm gonna pass on the temptation to talk about measuring GPUs and focus on what seems to be the theme of these three really interesting papers, which is sort of three feedback mechanisms that are gonna throttle the advance of AI.

One is the financial markets, another is the labor markets, and the third is the first paper, which is really this morality question, which I don't think we've talked about rigorously enough in the econ profession. And I think the comment I have for the paper is that there are actually examples in history.

You have to look beyond papal encyclicals. The Pope has come around to caring about these things relatively recently, but you can look back at the 60s at liberation theology after Vatican II. Those people were ostracized in the church, but they were out there talking about the effects of globalization on the Global South.

Look at what John Steinbeck had to say about agriculture. You can look at what Karl Marx had to say about the Industrial Revolution. And all these things are grounded in a sort of moral outrage. And I think the question for us is, you know, which of the outcomes from those historical examples are we gonna get, right?

Are we gonna get totalitarian oppression as we did in the Global South? Are we gonna get the New Deal? Are we gonna get the Russian Revolution or the Cultural Revolution? And I think it's really important for us to figure out even if it's really hard, how to predict where we're headed.

EBERLY: Katherine Abraham and then John Haltiwanger.

ABRAHAM: So I David has just made one of the comments I was gonna make, but I maybe just to add one thing to it. I really liked Martin's papers distinguishing among economic concerns, existential concerns, moral concerns. I think I would agree with David that it's not clear to me that this is really new with... I think we've seen similar concerns about past technological changes, but I wonder if maybe what's different this time is that it's us. It's people who do the kind of work we do as opposed to, you know, other kinds of work. But I'd be interested in your thoughts about that.

I did just have one question really about the last paper. I of course wholeheartedly agree about the need for better measurement. But but I do wonder with respect to part of this, for example, you know, Meta offloading its debt to these other entities, whether you think that investors are likely to be able to see through that. It's such an important part of what they're

doing, and I wonder, at least in that narrow piece of it, whether people might be able to see through what's going on

HALTIWANGER: So very interesting discussion, and I thought maybe David and Katherine would pre-empt me, but I'll go ahead and go into measurement issues even more. I think the measurement challenges here are enormous. The third paper particularly emphasizes this.

I also think, you know, in Luis' talk, which I very much enjoyed, I wanna caution against there's just a proliferation of, I'll say, private sector data sources that are very creative. By the way, some are surveys, for example, that the Richmond Fed or the Atlanta Fed is doing and down the line. If you look carefully at them, oftentimes these are small-scale surveys, highly non-representative surveys. So I would advocate in terms of Luis' point on diffusion, there's something called the Business Trends and Outlook Survey, which is a very large-scale survey by census, two hundred thousand firms every two weeks tracking all of this. And when you look at that, yes, AI is taking off, but particularly in some parts of the economy much more than others. It's in large firms and in particular sectors. So, I think that's really quite important. In lots of sectors like retail, small retailers, they're not using AI at this point.

Then in terms of the third paper I can't -- again, as Catherine said, we can't agree more in terms of trying to get the measurement right. And what struck me, here's where I'm thinking actually the private sector data sources might be much better. I'm not convinced we're doing a good job of capturing what's going on in the national accounts. I think we're probably getting-- I think the data's up to the task of getting the capital expenditures okay. But we don't stop there, right? 'Cause we wanna measure TFP, and so we gotta measure more capital stocks. So do we-- what do we need to do that? We need depreciation rates, and we need deflators. And I think we're in deep water on both of those. And actually, the third paper sort of hinted, yeah we're in quite deep water.

So, you know, people are already very much, you know, looking through the tea leaves of the TFP statistics to see is it showing up in the data yet, the AI. It's not clear we're getting it right, particularly on the data center side. We may be-- suppose depreciation rates are much higher for this kind of capital investment, and we're not capturing it, then we're probably understating TFP at this point rather than overstating. On the other hand I have no idea what's going on the deflator side. So I think the measurement problems are just huge here.

EBERLY: Thanks. Kim Clausing and then Joe Beaulieu.

CLAUSING: Thanks for those fascinating presentations. I learned a lot from all three of them. On the third one, I was really intrigued by the scale of some of what you documented there, but I wonder how sensitive both the fiscal and the sort of crowding out risk that we're thinking about on the investment side is to the international element of this.

So if we think about the enormous amount of build that has to happen in the United States and your calculations, like some of it can happen elsewhere too, and it partly depends on how this competition evolves. But it strikes me that both the cyclical and the crowding out implications are sensitive to how much will happen internationally.

The first paper I thought was really intriguing, and I liked the moral thing, although I'm not sure I would label it exactly that way. It seems like the loss of human dignity or jobs or existential risk is also intensely moral. Whereas the theft of IP almost seems more narrowly

economic to me. But as you were talking, it occurred to me that there were some really interesting parallels between this debate and the trade versus technology debate that I'm sure occupied this room for a decade or so earlier, where we wondered, well, inequality, is it caused by trade? Is it caused by technology? To some extent, the policy implications were the same regardless of what it was caused by, which is that you should find a way to redistribute, right, from the winners to the losers, because we don't wanna throw away technology, and we don't wanna throw away the gains from trade.

And it strikes me that something is really similar here. I'm not sure that the ownership solution that you propose as a response to the IP theft is really the optimal one. Like, if the government or society owns part of these companies, you know, it raises interesting governance questions in terms of the line between corporations and the state.

But it also, it's not clear that government and society will necessarily benefit if those firms ultimately fail for some of the fragility reasons that we talk about later. Whereas the tools that we already have for handling the redistribution, the dis- the disruption and the, you know, dislocation that this process will enable, we already have these tools with corporate profits tax will work regardless of whether the firms are profitable or not, or whether it's the doctors that, that take that, those benefits and, you know, big tax credits and safety net expansions at the bottom.

And those seem to work regardless of how this ultimately turns out without requiring the government to take a stake in firms.

EBERLY: Joe and then Bob Gordon.

BEAULIEU: So on the first paper, Martin's paper I think it's absolutely right that there seems that the trend in attitudes is being generated by something a lot more than just, "Oh, I don't want a natural gas generator in my backyard." That there's something much broader than that.

And I take the the identification, the idea of like what we can investigate is the same thing that it's always is, which is, well, what's different this time? And in that vein, I would submit that the much bigger thing that is different this time is the basic distrust of institutions in this country at least, but probably around the world. Gallup has polls that go back to like the early '80s in terms of attitudes towards various institutions. And the striking fact in that those poll results is that all, almost all institutions except for firefighters, I think, have been trending down. Some of them cratered earlier, like Congress, others a little later, like the Supreme Court. But all of them, except for firefighters, are way down.

And I would submit that's what's generating these things. So in the past, and, you know, Gallup wasn't running the polls when Vanderbilt was building his railroads. But I would submit that in the past it'd be like, "Well, see, the Pope's got it handled, so I'm not gonna, like, worry about this too much." But now if you just don't trust anybody, then any change at all is scary

EBERLY: Bob and then Josh Gotbaum.

GORDON: Oh, okay. Seems to me that we've got a sort of equilibrium that's likely to come out that I haven't heard about. Three point seven trillion in plans. At the same time, you've

got enormous growing political opposition to the data centers. And a lot of these three point seven trillion are not gonna get built because the local authorities, even states, are gonna start prohibiting them and slowing them down.

So we might get an equilibrium in between the overbuilding that is implied by what I've heard and the political opposition, so we could come out in the middle somewhere meaning less debt and less expenditure. The other thing that we wanna remember is that we heard that right at the end that revenue has to grow 80% a year for seven years to make this even plausible.

What about profits? Are profits even gonna be there given that we have growing open source free models that are gonna be available to people at very low cost coming from China, coming from all the venture capital startups that are happening all around us? So it's definitely an unstable system, but exactly which way it's going to go, I think is more complicated than three point seven trillion.

It's it's going to be less than that, and it's gonna work out, I think, in a different way, partly because of this open source competition.

EBERLY: Josh?

GOTBAUM: The three papers together actually struck me as almost elegant because yes, I think of this as a combination of a new technology, the economic risks of which are being passed around and diversified and expanded. And we've seen that before. We saw that in the 1980, in the 1980s, et cetera. And so one question is, yes, there will be a, there will be a reconciliation, and the question is, will the reconciliation destroy an entire class of assets or simply reduce returns on institutional investors overall?

At the same time, properly, you've said that there is a development of a rethinking of the property rights and et cetera. And rather than ownership, I would send you all the way back to the era of the railroads and the robber barons, which led to the rise of the national income tax. And so you're -- in effect, what you are doing is you're trying simultaneously to solve a risk diversification problem and an intellectual property problem, which is why I thought actually the combination of the papers was pretty nice.

EBERLY: Okay. Thank you. Let's go back to the authors then. So why don't you go in the same order that you spoke and I wanted to point out that Martin's co-author, Noam, is on the screen. Noam, you're like six by six in the screen in the front of the room. So, it's great to have you join us- And you can,

YUCHTMAN: Thanks for having me.

I don't know if I deserve to be so prominent but I'll- you know, make sure that I don't do anything too embarrassing.

EBERLY: So Martin and you two can share your response, and since we need to hear from three of you, try to focus your responses to the questions.

BERAJA: Yeah, I'll try to keep it short then. You just see, it's funny that some of the debates that we're having here are exactly the discussions that we have with Noam in, in, in private. And let me just touch on, on I think two.

One, so I really liked the first comments on how-- by Katharine and the one before, on how I think what is really different is, yeah, it's true that previous technologies have also raised moral questions and, but this time being about us, I think that's what makes it different.

And let me just add, I think one of the points that we had in the previous work with Noam was how in history, the way that these, you know, innovations intertwine with appropriation conflicts have been resolved, have often been, especially in the US, but have been resolved by legitimizing the taking exposed. And that's how, you know, the property rights institutions have emerged. And also institutions like, you know, the national income tax is a, is, or the New Deal, these are great examples.

And maybe, where we're coming down on this, at least myself, is that because this legitimizing the taking this time doesn't just mean legitimizing a taking of potential revenue streams potential economic implications, but it's really something else, and that has, you know, this human flavor to it. And so there's risks that have to do with what, you know, the first comment was about it. What, you know, what, what type of institutions, what societal change will we get if we really legitimize this this type of taking?

And I really like the analogy. I think, I thought when you mentioned trade and technology, I thought your comment was really gonna be about how in that context, the discussion was also about should we really compensate the losers with money or should... or is there something about giving jobs and that, that is what is more that people want jobs and that gives human dignity. And so by just giving a transfer, that's not gonna solve the problem of you know, human dignity. And and collective ownership, I think people interpreted collective ownership that meaning the government should take, you know, a stake on these companies and redistribute.

I don't think that's what Sam has in mind when he says collective ownership. I think that, the resolutions could be entirely private in this case. And I mean, this will be a much larger discussion. But it's really about recognizing the authors and all the creative work and and for, sorry, and for human expertise and workers, that will also be, you know, resolved internally within the, the companies that are using this data to train together with their worker unions probably. And so I don't think collective ownership means necessarily the government takes a stake on it. Yeah.

YUCHTMAN: Yeah, I'll just add one, one minor point, which is you know, all of this discussion relates to how we think a political equilibrium will arise out of what we anticipate to be meaningful conflict, and that was in the background of several of the comments that there will be some sort of, you know, potential rewriting of the social contract or debate over the social contract in a context in which there is a huge amount of political distrust in, in incumbent policymakers and in, in other institutions.

And so, you know, a lot of what Martin and I have been discussing is, you know, how can we think about a, a world where we also learn from the inability to effectively compensate the losers from, you know, international trade, for example. And the aim to actually give people who feel this grievance a sense that grievance has been addressed in a meaningful way.

We think that something like ownership would be, you know, one, one, you know, step forward toward thinking about addressing citizens' needs for something like dignity. Recognizing their sense of grievance that something, again, very personal seems to have been taken from them, from us.

And thinking, you know, down, down the road what that means about politics. Right. So, so your points are all really well taken. Great. So thank you.

EBERLY: Thank you. Luis, do you wanna respond?

GARICANO: Yes. Thank you. Great questions and comments and, so on measurement, Arvind I think that what you wanna measure is what is the minimum cost of acquiring a capability, and how is that changing? To what extent do we see Microsoft doing its own corporate law and that kind of thing? I would think that you can do it experimentally. You can try to do that experimentally, give work to a team and give work to individuals. There is a P&G experiment by, in Procter & Gamble by Rafaela Sadun and a team of co-authors doing something along these lines.

You could also measure the amount of problems that get passed internally on inside the organizations. How much, how many times do you call the IT relative to how much did you used to call it before? And you could do it between organizations. I loved this, the point about the slowdown of technology being common to the three papers.

My sense from my last two or three slides that I couldn't really go deep into is that The heterogeneity in growth rates over the next ten years is going to be driven fundamentally by the political economy. That to the extent that the lawyers and the doctors and the truck drivers and the bus drivers organized and stopped the adoption of technology in some countries and not in others, you will see very sharply different returns on Stein's investments.

And I think that one aspect of what Martin was saying that is critical is that the ones who are affected now are not the disenfranchised people in society. They're the powerful people in society. They're people like us. They're the lawyers and the doctors, and that was-- that's kind of the groups that my paper is talking about. And those people have the ability... I mean, in Europe, twenty thousand notaries are stopping the effort that European Union is doing to incorporate firms at the Euro-wide level because they all want to be able to look at the documents in each country. Imagine if twenty thousand notaries can do that, imagine what all the lawyers in Europe can do together.

So, so on the, on, on Jon's, on Jon Haltiwanger's point to be careful with small cases. I take his comment. It's true. There's a lot of data sources. I loved the, the ChatGPT guys because they actually get at how much am I asking outside of my domain. I feel myself asking a lot of stuff that is not normally my domain. So, I kind of see that as a trend that is probably easy to - possible to identify this, this referral probability. But I will ask him for this census data, Jon, and we talk later about what can we do with it to get closer to Arvind's question.

Which, by the way, just the last point, it has difficulty, which is that I'm arguing that there are two things going on together: the drop in transaction cost and the gain in the knowledge cost, and that they're racing against each other. So in order to see it in the wild, we're gonna have to be able to think of those two empirical objects. How much easier it is to ask the

specialist because transactions are cheaper because of AI. And I just have an agent who goes around. And how much less valuable it is because knowledge is cheaper.

And I want to finish by asking Stijn. I know you have too much, so maybe you can ask my question. In Europe, we are six percent of that of that build-out. We don't have data centers, full stop. I mean, now we have a big advantage because we don't have data centers, we don't have politics either. So people are not as angry as they are in the US. So we maybe have a little window to do this, but we don't have any demand guarantee, and nobody guaranteeing the offtake because the little firms in Europe are not willing to do this. So I wonder to what extent and how would your financial engineering help us? But maybe I should ask this offline.

EBERLY: Okay. Let's turn it over to Stijn. We're gonna run a little late given the gravity of the discussion we're having. So we wanna hear from Stijn.

VAN NIEUWERBURGH: Thank you so much for all the great questions.

Let me start with Arvind's question. It's a good one. So I think that if we moved all of the-- So first thing, the first point to make is that the hyperscalers are a big part of the story, but they're not all of the story. They're maybe half of it roughly. But I think you're right. Ultimately, the leases to third-party landlords is what's backing all of this build out. And these leases are pari passu with corporate debt, so they're very severe obligations. Meta cannot walk away from these leases easily. It would have to go bankrupt to walk away from it. So ultimately, what we're talking about is bankruptcy risk of one of these large hyperscalers.

Now, as unlikely as that seems, given their credit ratings today, we have to remember that they're in an arms race with each other, that they're all betting the house simultaneously, and that they probably cannot all win this arms race. That said, it's a remote possibility today that one of these four or five, you know, three, four, \$5 trillion companies will go bankrupt. I think if we were to move all of, let's say, the \$10 trillion build-out is financed with 70% debt, \$7 trillion of debt, if we were to put that on their balance sheet, their leverage might be around 45%, would be, you know, a bit higher than your average corporation, but seems still manageable.

I think the risk is that one of them will go bankrupt. The other half of the ecosystem is much more fragile. Now we're talking about much smaller companies, much less well-funded companies like CoreWeave. You know, the CDS spread on CoreWeave is 10% per year, right? So there's clearly a lot more bankruptcy risk there. And so I think that's, you know, we should keep that in mind.

I think you're right. It's interesting that there doesn't seem to be a lot of short-term debt in this system. Sort of the only thing that approaches this a little bit is these GPU-backed securities. These are variable rate, and they tend to be relatively short-term, given that the life of the GPU is maybe four or five years. But it's not asset-backed commercial paper, right?

Katherine you asked about do investors see through this? It's a good question. I think it's a mixed bag. If you look at private price earnings ratios of some of these companies, you would think that investors have you know, dollar signs in their eyes and it's pie in the sky. If you look at some of the corporate bond yields and the CDS spreads that's maybe a different group of investors. They seem to price some of this in already. So corporate, a lot of these

companies, including the hyperscalers, they're investment-grade credit, yet the bond yields that they're paying are more like junk bonds. So there's an extra premium that bond investors are already demanding for holding their paper, and the same in the CDS spread. So, maybe two different sets of investors are thinking about this risk differently today.

I like both Kim and Luis's question about the international dimension of this. A lot of these same US hyperscalers are actually building data centers all over the world. And you could sort of imagine with all the political resistance that they're meeting in the United States, that it might become easier to go build them in the Philippines or or somewhere else. And that might, in fact, for model training, be just fine because those data centers don't need the low latency of inference. So I think we might see more of that in the future.

It's still, I think, you know, gonna get funded largely through global capital markets and, you know, predominantly, I would say US investors that are funding all of this. So in terms of the crowd out, I think it's a complicated question because there are global investors and there are global investments on both sides.

I think, Luis, you're right. I mean, Europe has a problem because the data center build-out is being done by China and the United States. There's very little happening in Europe. There are very few AI companies and there's very little data center capacity. Some, because of data sovereignty, some computations have to be done locally. And then there's a question as AI adoption grows, it seems to me like Europe might very quickly run into a binding compute constraint unless there is, you know, a real concerted effort to bring more capacity online.

And I think, you know, Mario Draghi pointed this out recently that some of it has to do with the fragmentation of the demand for compute. There's a lot of little companies maybe by forming some sort of, you know, collective, they can try to solve that problem through contracting that way. I think there might also be an interesting question about the, the capital markets. Are the capital markets in Europe deep enough to finance some of this?

Bob, I thought about this possibility as well about, you know, maybe we won't be overbuilding because of the political opposition. It's sort of an appealing middle ground. I think in the short run, it seems like some of these projects are being stopped in their tracks. And by the way, remember my calculation threw away two-thirds of the pipeline of announced projects. So I'm only considering one-third already reflective of the fact that there is a lot of opposition and a lot of sand in the gears.

That said, remember the timing is important. If you have already issued the debt and you know, you know, the data center will be finished three, four years from now, you're already on the hook for servicing this debt. And if you cannot build this data center, you have no revenue to pay that debt. So that is still a problem.

I also agree with your point on revenue pressure. I think there is a lot of competition. I mean, I tried to make this point that how likely is it that we can quadruple our data center capacity without any reduction in pricing power, right? There's all this competition from Chinese models, from open source models. It seems very difficult to charge these frontier prices on five times more capacity, yet that's what it takes to you know, to generate the revenue to generate the return on these investments. So I think that is sort of a tension point.

And then finally, Joe sort of made an interesting point about diversification, right? One thing about the great financial crisis we initially thought was it's wonderful that we have all the securitization because it's spreading the risks away from the levered financial sector towards unlevered investors, mutual funds, pension funds, sovereign wealth funds. Those are the people who can take the losses. Those were the people that took the losses in the internet bubble, by the way. It was rich people paid for it. The Nasdaq fell 80% and rich people lost money in their 401k. And that's bad and it has consumption wealth effects and all that, but it's not a financial crisis because it doesn't run-- it doesn't impair the banking system.

During the GFC, we thought, you know, all this housing, all this subprime mortgage is fine. It's all gonna get dispersed throughout the system. But then later on, we learned that a lot of the risks were taken back through the back door on banks' balance sheets. First loss pieces were held on by the banks, and then the banks were giving guarantees and the insurance companies were giving guarantees on some of these tranches.

Some of this is happening today again. So yes, we have securitization, but we have guarantees. Insurance companies are selling guarantees on some of these tranches again. We have other things that really are eerily reminiscent of what we went through. For example, data center mortgage-backed securities are exempt from reporting requirements, and they're exempt from risk retention rules under Dodd-Frank. Why? And I point out several of these things in the paper that sort of have an echo to what we went through. So yes, if we can spread the risk away from the banks, great. But by my calculations, the banks today already have at least seven hundred billion dollars in debt exposure to data centers and are up against their concentration limits.

EBERLY: Great. Thanks very much on that upbeat note. So we'll we're gonna extend the break a little bit, come back at 10:35, and next time Josh will aspire to be actually elegant. Okay.