ECONOMIC INNOVATION GROUP

THE NEW BAZAAR

MAY 16, 2025

AI and Jobs: What Do We Really Know?

NATHAN GOLDSCHLAG ON HOW BUSINESSES AND WORKERS ARE USING THE TECH OF THE FUTURE

CARDIFF GARCIA: Hey, everyone. Cardiff here. Artificial intelligence, AI: is it gonna help you do your job better so that you can produce more of whatever it is that you produce? Make a lot of money, have a happy and fulfilling life with your brilliant new AI assistant?

Or rather than help you do your job, is AI just gonna straight up do your job? No need anymore for those of us who do knowledge work, laptop jobs, thinky work. Because we just will not be able to compete anymore, and there goes our usefulness to society, our employability, our self-worth, our dignity.

Or will this all amount to something else? Something in between those two extreme examples, or maybe something that we simply cannot anticipate now.

Today's guest is Nathan Goldschlag. Nathan is an economist, and until recently, he was at the U.S. Census Bureau's Center for Economic Studies, where his title was Principal Economist, and where, among other things, he led research on the impact of technology, including AI, on the economy.

And by the way, any list of the world's best economists on this topic has gotta include him or else your list is totally worthless. Sorry. And happily, he also was an easy guest to book because he is now the Research Director at the Economic Innovation Group, which happens to make this podcast, and also happens to be where I work. So he's now my colleague.

And what I've learned working alongside him is that he has just a profoundly deep understanding of the available data on the effects of AI on the economy. And you can see it in his own work, his own research, but it also makes him uniquely qualified to filter out a lot of the nonsense out there about this topic — and believe me, there is a

lot of nonsense out there about this topic — and qualified therefore to find the outside analysis that really is worth paying attention to.

So what we do in today's episode is we not only go through Nathan's own research, his findings about AI and work, we also go through some of the work done by others — researchers, essayists, Substackers — and we go through all of their writings and analysis and meanderings that we think is worth sharing with our listeners.

And of course we'll also link to all that work on our website at eig.org. I couldn't wait to get Nathan on the show for this chat. Here it is.

Nathan, your debut on The New Bazaar. How are you? How you feeling?

NATHAN GOLDSCHLAG: This is great. I'm super excited to be here. I think some of the work I've done at Census gives me a good foundation to think about these questions, so I'm super excited to share some of the things that I've learned along the way.

CARDIFF: Yeah, it sure does. And also because you're not at Census anymore, you can really say what's on your mind.

NATHAN: (CHUCKLES)

CARDIFF: I like talking to people who just left the government!

NATHAN: That's right. That's right. The cuffs are off

CARDIFF: Unshackled. I love it. Do you think a lot about AI and its effects on your own job?

NATHAN: Yeah. I mean, well, in some sense I think that the current iteration of LLMs are super useful in the work that I do. I don't consider it to be able to do the work I'm — since EIG just hired me.

CARDIFF: LLMs like ChatGPT, the large language models that you basically ask a question and they tell you anything.

NATHAN: Right. It can't do my job, thankfully.

CARDIFF: Yet? (CHUCKLES)

NATHAN: Well, yet. Maybe never. Hopefully never. And so I think those technologies right now they're super useful as an aid to brainstorming.

CARDIFF: Aid to brainstorming. Yeah. I think about this all the time in the context of my own job. And when I look at you and what you do, you do a lot of things.

One of them is you manage people. LLMs, ChatGPT can't do that, right? You oversee people, you sort of have to use your judgment and your accumulated wisdom from the years in which you've been an economist and manage other people.

You prioritize what it is that the researchers and economists who work for you are gonna do. So there's a managerial component to it, but you also are an economist yourself. You do a lot of original research of your own. And how does AI help with that? So it helps you brainstorm. Does it help you do things like the coding? Does it help you, I don't know, check to see if you've made mistakes? Things along—

NATHAN: Yeah, absolutely. So it's really helpful in coding. You don't have to hold all of that stuff in your head anymore. You can just let it go and let the machine remember for you. That happens a lot. You sort of say, "Oh, how do I do this task?"

And you can just ask it and it'll give you a bunch of pieces that you have to string together oftentimes. For the work that I do, I think that it's all about research taste, right? And so it's hard for these systems to have the ability to know what questions to ask, and that's what makes it really useful for a brainstorming tool because you can bounce ideas off of it. And you see the things that it's giving back and can string the more important and interesting pieces together. And that's the value-add that I thankfully still have.

CARDIFF: I know a lot of people who use LLMs as an aid in the coding part of their jobs, but they all also tell me that they have to check the code that ChatGPT and the LLMs give back to them. So it can help you save a lot of time if you say, "Hey, give me this code. I want to do the following things." And I'm sure that these tools will get way better, but then they do say like, you have to do, you have to be very careful after you've made that request, and it gives you back the code and it has saved you the time. You have to then spend some time making sure that it's right. Is that your experience too?

NATHAN: That's right. It just gets you a lot of the way there, right? It'll give you something that still has some bugs, and it forgot to do some things, but again, it saves you a huge amount of time, and it gets you most of the way there with the code that you're trying to write with way less effort.

CARDIFF: Yeah. Some of the new AIs are also doing things like releasing podcasts with fake AI podcast hosts that actually talk to each other.

NATHAN: (CHUCKLES)

CARDIFF: I think Google released a new program that does this last year. I think about AI in my own job all the time. I sort of have a feeling that a lot of the individual subcomponents of my current work will be doable by an AI within, I don't know, five or ten years.

Like the talking part, clearly. What you and I are doing right now on this podcast. Like it's already started.

And in terms of helping people write better, 'cause I do a lot of editing. I do a lot of writing. I don't know. It seems like a lot of that actually is automatable, and at the rate at which these tools are getting better, because they are improving all the time, it seems like a lot of the parts of my current job will be doable by these tools later on.

And so I will personally have to evolve what it is that I bring to the job, as well as maybe use these tools to make me better, and not just keep doing things the old way. Do you think, as an economist who studies work and labor, do you think that's good advice?

NATHAN: No, I think that's right. And I think that in the future, to be successful, you're gonna have to be leveraging these tools, and in order to make them make you better at what you do. Any one piece of the things that you do, yeah, that could be automatable. But the whole thing, you gotta remember, there's an agency problem of stringing a bunch of different things together. And so each individual task taking on its own, maybe the AI can do that piece. It could also help you do that piece, so it's a compliment as well as a substitute. But stringing all those things together is really, I think, the value-add that you're gonna retain.

But at the same time, the tasks that you do are gonna change. What tasks you focus your time on, and the allocation of your time across tasks, is gonna change. And it might create new tasks as well. So you might be doing things that you didn't do before because these tools are opening up new doors.

CARDIFF: Do you think that raw intelligence will no longer be as valuable as it once was? Leave aside the issue of taste, which you mentioned. I like that point. The issue of judgment, wisdom, what to do with the output of raw intelligence, just like sheer cognitive power. The ability to remember stuff, the ability to make certain

connections. It just seems like with the elevation of these tools, that's just not gonna be as impressive as it has been for the last few decades.

NATHAN: No, that's right. But it's also the case that in the data, the returns to IQ for wages aren't nearly as high as you might have thought. So it's not as if the labor market is actually—

CARDIFF: Oh, that's why I've never made much money!

NATHAN: (LAUGHS) Yeah, no, it's all the social skills.

CARDIFF: That's right.

So, in some sense, people are overrating IQ to begin with. But it's also the case that, as the supply of intelligence increases, right? There's gonna be diminishing returns to that. So there's gonna be complementary skills and assets that humans are bringing to the table that make that equation work.

CARDIFF: Yeah, better to pivot to those as soon as possible.

NATHAN: Right. Right.

CARDIFF: Let's talk about your research, 'cause you've published a number of research papers about AI and technology generally, and their effects on labor market. Let me start with this. What do we know about how many, or the percentage of businesses, that are actually using AI as of right now or the most recent available data?

NATHAN: So this is actually something that I've been spending a lot of time on, and a lot of the work that I had done at Census was focused on these questions. And there's a really great team of economists at Census that are thinking ahead of the curve on these things.

Oftentimes for the federal statistical system, the technology is fully diffused by the time they ask the questions. Here, in this case, we actually started asking them early. So we actually have a picture of what this looks like over time, and I think the bottom line is it's lower than you might've expected.

So in the most recent data, something like 9 percent of firms say that they're using AI. Pre-GPT, that was maybe 3 percent. So it's increasing over time. But if you ask businesses, "Do you use AI in the production of goods and services?," only about 9 percent of them say yes.

CARDIFF: Yeah. That's really interesting because that does sound low given the prominence of the debate about AI, and the culture and the media. It's obviously something that you and I are very interested in. That's lower than people think. Do you think that it's going to increase quickly though? How does this compare to before the introduction of those LLMs, these large language models, ChatGPT and so forth? Because that was just a couple of years ago, so I wonder if now the uptake's gonna be faster than it was before?

NATHAN: Yeah. The great thing about some of the work that Census is doing is that it's on these really high frequency surveys. So the data that I'm referencing is actually from May, and so these are really new—

CARDIFF: May of this year?

NATHAN: Yeah, this year. Right. So just recently.

And so these data show a trend. They're increasing, they've gone up from something like 4 percent a year and a half ago to 9 percent today. It's also the case that the firms that adopt these technologies tend to be much larger. So, early on, a couple years ago, that 3 percent use rate turns into something like 12 percent on an employment weighted basis. Meaning the number of employees that are at businesses using these technologies, right?

CARDIFF: Ah, that’s interesting.

NATHAN: They tend to be much larger. So the quote unquote exposure to them is higher than the firm level use rate.

CARDIFF: That's really interesting. I also imagine that different kinds of businesses show a lot of variation in how much they use it, right?

NATHAN: That's right. There's a strong relationship between size and the use of these technologies. It tends to be larger firms that use them, but it also tends to be younger firms. So within different size groups, there's a negative relationship between use and age. So it tends to be larger younger firms.

And one of the things that we've seen more recently in the post-GPT period is sort of a J shape in the size distribution, meaning that the very smallest firms have a little bit higher use rate than sort of the mid-sized, and then it goes back up. And the largest firms have a very high rate. And, we think that might be suggestive of the fact that these new versions of AI are really amenable to the production technologies of small firms.

If you have a firm of only a couple people, you might be able to use these technologies. You didn't actually hire a marketing person before. Now you can just use ChatGPT to do something that looks a lot like a marketing person, right? So there might be higher returns among the smallest firms.

CARDIFF: That's really interesting. I would also imagine that if you're a small business and you're making the decision of whether to start hiring more people, increase the scale, you'd be careful with that decision. Because if you're a small business, hiring another person is a really big prospect. So you might just be more experimental with these technologies first. See if you can save the money of not going from five people to six people, which is, that's 20 percent increase in the size of your business versus a really large company, which might just be like, "Hey, we have some more work to do, just hire somebody else to do it." You'd be maybe more willing to test out this technology if it might save you money if you're small.

NATHAN: That's right. That's right. That sort of feeds into the what are the employment effects of these technologies.

One of the nice things about these surveys is we can ask firms, you say, "Did you increase or decrease your employment based on your use of AI?" Like it's a self-reported causal estimate, and when we do that, most firms, the overwhelming majority of firms, the ones that use AI, say that AI did not have an impact on their employment.

And in terms of increase and decrease, it's about 50-50. Like, among the small percentage that say that AI did change their employment, half say up and half say down.

CARDIFF: I would also imagine that some companies would wanna hire people to help them use AI. So it's not just that AI might automate away the work of some employees. There might also be some employees you'd bring on who happen to be experts in the technology, and you'd bring them on as well, right?

NATHAN: That's right. And we have estimates of the skill composition. So we ask firms whether or not they're hiring more skilled workers because they use AI. And the answer tends to be yes, more than no. And so it's also true that when firms are implementing these technologies, it increases their demand for skill. They need new types of skills, and they need new skilled workers to implement those technologies and reap the benefits.

CARDIFF: Yeah, that's interesting. And I can understand, as we just said, why the smallest businesses would be willing to experiment with AI. Bigger businesses may

not be as thoughtful about trying out at the very beginning, the thing. But the very biggest companies you said also tend to use it a lot more.

And I wonder if, in their case, it's just because they have the money to invest, to upgrade their systems or whatever, to accommodate AI. Because they need, some companies do need to change the way they do things to use a new technology. Isn't that right?

NATHAN: No, that's right. And there's certainly gonna be returns to scale. One of the things that you see this in is with data. So these large businesses, they tend to have lots of data. And that data you can use AI to increase your returns of the adoption based upon that data. You can sort of customize these tools and implement them in specialized ways that leverage the information within the company.

CARDIFF: Yeah, I like that. And then if companies aren't using AI to automate away worker tasks, or maybe some are, some aren't, but it doesn't look like it's super prevalent just yet, what do they say they are using it for?

NATHAN: So the most common response is to improve products, and improve processes. The third most common response is to automate the work done by employees. But among those, among firms that say that they're using AI, about 27 percent say that they're using it to automate tasks, but only 5 percent say that they're changing their employment. So it's more often the case that firms are using AI to automate subsets of the things that different types of workers do, not necessarily do all of the things that those workers did.

CARDIFF: Big difference between, let's say a company that mainly does, I don't know, internet stuff, information technology also, versus companies that maybe do more stuff, I don't know, in the physical world, in the tangible world. I would imagine that the adoption rate of AI there also varies quite a bit.

NATHAN: Oh, it varies a lot by sector, so it can get as high as 25 percent in the information sector, but something like 2 percent in the accommodation, in food services sector. So, one of the most common responses to why a firm isn't using AI is that it's not applicable. A lot of businesses still don't see how it can improve their production.

CARDIFF: Did anything about all this research surprise you as you were doing it? Where you put out the surveys or you looked at the data, you and your colleagues, and you saw it, and then you were just like, “That's interesting. Didn't know that was coming.”

NATHAN: So I think the overall low use rates were a surprise. There was like within a week ChatGPT had a hundred million users or something. There's amazing statistics about diffusion.

CARDIFF: Where do they work? Where are they? (CHUCKLES)

NATHAN: Yeah. And so, based on the headlines, you'd think that it's sort of everywhere.

And based upon the way that economists are talking about it and sort of modeling out the relationship between AI and labor markets, you'd also expect to see lots of displacement. Lots of movement employment among adopting firms, and that just doesn't seem to be the case yet. And we're well into the post-GPT period, right?

These technologies have been around for a little bit now, but even in the latest data from earlier this year, it doesn't look like these companies are changing their employment a lot, even though they are expected to. Something like 35 percent of them, in the next six months, are gonna be using it to automate tasks.

So I think the automation of specific types of tasks that workers are doing is gonna go up. But I think if you, taken on its face, these low percentages of use and the not rapid displacement of workers is kind of surprising.

CARDIFF: This might be too granular for the data that's available, but I was talking to a friend of mine who's an executive at an investment firm recently, and one of the things he was saying was that they haven't figured out how to use AI yet for the main thing that his firm does.

What they use AI for are things like customer service or investor relations. These more kind of peripheral tasks where yes, they can automate away some of the work that the workers do, or they can use AI to enhance the workers in those parts of the company. But it was only now that they were starting to figure out ways to use AI in the main thing they do, which is essentially financial analysis and investing.

And I'm wondering also the extent to which a lot of the companies that respond to the survey are also saying, “Yeah, we use AI, but right now at least, it's for things that are vital, they're important to what we do, but they're not the main thing that we produce.” You know what I mean?

NATHAN: No, I think that's actually points to — it's an interesting point, and I think it actually points to one of the challenges in measuring the use of AI among businesses.

And it's actually about whether or not — if you think about our low percentage of uptake of AI, that might reflect the fact that asking them whether or not they use it in the production of goods and services is kind of a high bar. There might be lots of incidental use within the firm.

Some of it might be embedded in products and services that it purchases. Cloud services or your email has lots of AI baked into it, but you wouldn't consider that part of your production of goods and services necessarily. You might be using it to summarize different pieces of information, memos coming in or going out, summarizing email chains or whatever.

Those incidental uses might not be picked up in these percentages. So, it's true the diffusion might actually be much higher, but the specific type of diffusion we're looking at of the production of goods and services, which by the way is where we think the most TFP changes are gonna be coming from.

CARDIFF: TFP. That's Total Factor Productivity.

NATHAN: That's right.

CARDIFF: So that's like the most essentially technological change that helps a company become more efficient, innovative, better, more productive.

NATHAN: Yeah. When we think about the implementations of AI that are gonna reap the most economic benefits and make us better at producing things, it's gonna be in the production of goods and services. It's not necessarily gonna be that the auto complete in your email can be exactly right.

CARDIFF: (LAUGHS) Much though I appreciate it.

NATHAN: The perfect emojis at the end of every sentence for everybody. That's not necessarily gonna move the needle on productivity.

CARDIFF: You found something really interesting that you mentioned to me in a recent conversation about the difference between software developers and programmers. So give us sort of the context here in which you were looking at this, and then tell us what you found and what's significant about it.

NATHAN: So this actually feeds back into the discussion earlier about tasks versus an agency, and thinking about the role that we have in relation to these technologies. So one of the often cited statistics when we think about the impact of AI and labor markets is the displacement of software programmers.

And one of the really interesting things that we can see in the data from the Bureau of Labor Statistics, where we track the number of software programmers, is that it's actually split between changes in the software programmers and then software developers. So software developers, who are much more numerous, they're not declining nearly as much as software programmers.

CARDIFF: Tell us, first, what the difference is between software developers and software.

NATHAN: So the software programmer, their primary tasks are related to creating, executing, and testing new scripts and programs. Whereas the software developer is much more deeply engaged in the why the program is being written and the requirements surrounding it, and stringing together multiple pieces of the business technologies to solve these questions.

CARDIFF: Is it fair to say that what software programmers do is a little bit more routine oriented than what software—

NATHAN: That's exactly right.

CARDIFF: Okay. And so when you had this finding that software programmers, in percentage terms, are suffering in the labor market more than software developers. There's a possible reason for it, but you also found that there's a lot fewer software programmers in the first place than software developers. Is that right?

NATHAN: That's right. Software programmers between 2013 and 2014 were down something like, in counts, they were down something like 8-9 percent.

CARDIFF: Okay. It's a lot, by the way.

NATHAN: It is, that's nothing to sneeze at. And the software developers was down less than a quarter of a point.

CARDIFF: Oh so, basically flat.

NATHAN: But at the same time, software programmers is only something like 120,000 of them in 2023. But there's about 1.6 million software developers. So when we're thinking about like these stories about the displacement of software programmers and developers, you really gotta split the different types of occupations by what types of tasks they're doing.

And it's absolutely true that the more routine tasks, like writing and testing scripts, are gonna be much more affected by the LLMs we have right now than something that engages more with business processes and thinking about why programs are being written

CARDIFF: Something that's been really interesting about this whole conversation to me, and which I've brought up on the show before, is that it was only around like the mid-2010s, that this famous phrase started going around: learn to code. Everybody was supposed to 'learn to code' because that's the job of the future.

And now there's a lot of people saying, "Well, wait a minute, coding is one of the tasks that's most vulnerable to now being done by artificial intelligence." And it's almost like you get whiplash from this because the responsible thing 10 years ago was to learn to code. And less than a decade, or about a decade later, "don't learn to code, learn to do something else," is now what everybody's saying because of the rise of these AIs.

And I think what you're finding about developers versus programmers does is it adds a little bit of nuance to that story, which is that there's still coding-related tasks, coding-related jobs that you can do even if the LLMs themselves are now actually writing a lot of the code.

NATHAN: Yeah, so that would've been terrible advice if you to take…

CARDIFF: (LAUGHS)

NATHAN: …if you were gonna take the very narrow version of that and just learn to write and execute scripts.

CARDIFF: Right.

NATHAN: But if, at the same time yo u learned how to think about business analysis and the writing of requirements for these different programs, thinking about how it integrates and solves problems for the business, like those types of skills, which again, more broadly when you think about how we're gonna integrate with these technologies, you should be thinking about broadening the tasks and changing composition of the tasks that we're doing.

And so this is sort of a nice microcosm of that. So if you're super narrowly focused and you're just writing programs, that's gonna be really susceptible. But if you're engaged in these complimentary tasks, tasks that are very highly agentic, meaning

you've gotta integrate lots of different pieces and lots of logical pieces of a chain, those are gonna be less susceptible.

CARDIFF: There have been a lot of articles written recently about how younger people are not getting hired at maybe the same rate that you would've expected them to, specifically college-educated younger people in their early or mid-twenties. And that the unemployment rate for these folks is now a little bit higher than the unemployment rate for the whole economy, although it's quite low in both cases.

And again, I'm talking only about recent college grads. Do you make anything of that? Like, is it possible that that's related to the fact that a lot of businesses are trying out these new technologies, at least to replace some of the more entry-level work at, I don't know, banks or law firms where they would've previously hired these high-skilled young folks to come in and do that work?

NATHAN: It's a super interesting trend. I don't know that I have a very clear answer for you but I will say that it started to turn actually after the Great Financial Crisis, right? So, when you think about the decline in that ratio of the all employment rate to the recent grad employment rate, it turns really, really early, in terms of the AI timelines. The types of AI that we had in the 2010s were not nearly as capable as the ones that we have now, right?

And so it's not as if the decline in that rate was reflecting those specific types of technologies. Now, I think there's probably really complex structural reasons for this. There's some fun things you can think about, like whether or not the signaling value of education is declining.

The elite overproduction hypothesis: we just have too many of these people. You can think about skill mismatch between the types of skills that recent grads are coming out with and the skills that firms demand there. There's a ton of different explanations, but I don't know that AI really fits the bill because it doesn't fit the time series trends.

CARDIFF: I was listening to a podcast that Tyler Cowen, an economist that you and I both follow, and he was talking to another AI expert. One of the questions that Tyler had was, “What's the worst age to be right now as AI keeps getting better and better, in terms of just being able to incorporate it into your life and into your work?”

Tyler himself suggested that the worst age was 40 because you're along in your career enough that you might think that retraining is too much to take on, but you have enough of your career left that it actually would be very useful for you to be able to adopt this new technology, and use it in your day-to-day life.

But I'm kind of curious, Nathan, to know what you think about that question.

NATHAN: Yeah, I think, I would say there's actually maybe a dumbbell pattern to this. 40 might be really bad if that that age group is less amenable to changes in their work patterns in reflection of these technologies.

Everybody's gonna have a tough time because it's just moving so fast. So that's one thing to keep in mind. But another age that might not be great is my kids, right? So they're gonna be entering a school system that's gonna be thinking about adopting these technologies. And it's not necessarily well thought out how that's gonna work in the educational system.

So, being really young right now and being exposed to this, what's probably gonna end up being a giant social experiment of including AI technologies in the public education system. And then also being in the workforce later in your career, and less able to—

CARDIFF: You'll be in an education system that's trying to keep up in real time with how it works and might not be able to.

NATHAN: Right.

CARDIFF: Right. Yeah. Fascinating. So it's like people in their early forties and five year olds.

NATHAN: Yeah. (LAUGHS) Those are tough spots.

CARDIFF: Everybody else is gonna be okay. I wanna do long-term speculation here with you because we've been talking about what we know to this point, but you are open to this story changing, right?

Are there things we should be following over time to see if, for example, the capabilities of these AIs go from complimentary technologies to people with good taste and with good judgment and it takes away some of the sort of raw intelligence or the routine tasks to something that exists more in the creative realm.

We have to be open to that story, right? We should be open to many varied futures. And I'm curious to know what should we be following to see if there is a big change in the trajectory of the adoption of AI and its effects on the labor market?

NATHAN: There's a couple things I would be looking for. I'd be looking for changes in those use rates at the firm level. Those are gonna be important and they might not

tell you the whole story as we discussed. They might only be capturing a certain type of implementation of these technologies within firms.

Person-level use rates, I think that these sort of studies that ask people whether or not they're using the technology, that's gonna be important. Labor force participation rates and shifting occupational employment distributions, right? So, the Bureau of Labor Statistics puts out these really nice matrices that tell you how many people are working in different occupations over time.

And so you can string those together and see how the occupation mix is responding to the implementation of these technologies. And I think the one takeaway I would suggest you take from this is that it's very highly uncertain, right? It's very unclear how this is gonna impact the trajectory of these innovations of labor markets going forward.

And part of that's gonna reflect the fact that the implementation of these technologies may create new tasks, it may create new jobs, right? And this is sort of the history of technological change over centuries is the — once these big general purpose technologies hit, they totally reorder society and labor markets, right?

And so they create new jobs that you wouldn't necessarily have thought would've existed before.

CARDIFF: A couple years ago, I was talking to an economist named Avi Goldfarb on this podcast, on The New Bazaar, and I asked him about this and I said, "Hey man, let's say that my specific tasks are in fact automated away." So, things like talking on a podcast knowledgeably about the economy, being able to assess a lot of information at once, some management skills, things like that.

"But I am left with one still useful quality, or a few maybe, which is just people skills. Just charisma, the ability to make people laugh, whatever. Will that save me?" And he said, "Well, maybe. But the thing to remember is that a lot of other people who do what you do now will also suddenly be competing in that same space. So you're gonna have a lot of other people who are also automated out of their jobs, and so you think you've got this angle because you've got this one thing that hasn't been automated away, but you also have a lot more competition from other people who are now gonna also be trying to do the same thing that you're doing because precisely for the same reason."

And I thought, well that's disheartening, isn't it?

NATHAN: Yeah, that's right. And I think it's an interesting point that the returns to different types of skills are gonna change in response to these technologies.

But again, I would point to the fact that new jobs and new tasks have been created throughout the history of technological change. And so, in some sense, that's gonna be the thing that saves the day is that as the supply of social labor increases, the set of ways that firms can engage with that is gonna increase as well.

CARDIFF: It might also just be super random stuff that we haven't imagined yet, precisely because we anticipate it. I sometimes think about chess because as of three decades ago, the machines still couldn't beat humans. And now there's no human that can beat even a mid-tier chess machine.

And yet, people still are interested in human versus human chess. Precisely because they don't wanna see a machine do that. They want to see two human beings competing in a certain way, similar to with sports. You could come up with a machine that never missed 30 feet away if you're shooting baskets or whatever.

But we watch great athletes compete against each other because we want to see the human element itself. But there's something about that that's also more about our entertainment, our amusement, our presenting ourselves to each other. It's a little bit different from what the labor market looks like right now and has looked like where we're sort of producing things that we buy and sell. Instead, we'd be buying and selling our ability to capture each other's attention.

And when I look around, I sort of see what looks like maybe the beginnings of that. I'm being very speculative here, to be clear, but you see this in influencer stuff, you see it in the survival of chess itself and in other places. And I don't have anything in particular to ask about this. I just find it interesting. I'm curious to know what you think.

NATHAN: Everything's sort of downstream of culture, right? So the implementation of these technologies is gonna be within a social and a cultural context. And it makes perfect sense to me that people are gonna put greater value on human generated content, right? And so, it might be the case that the AIs can make music or whatever it is, but I might still pay more in order to know that there's a human on the other side of that engagement.

CARDIFF: I now want to do something with you where we go around the horn of some other research that other scholars and other writers have done. Because you gave me this list that's really interesting of stuff that you found worth engaging with.

Doesn't mean you agree with everything on this list, but you think this is all very much worth engaging with.

So we're just gonna go through them one by one and then we're gonna share it with our listeners in the show notes. So, first up, some economists who go by the names Bick, Blandin, and Deming, they found something that on the surface at least appears to contradict your earlier findings about the low usage rates at businesses.

But you have a way to reconcile them, and I think a lot of people are gonna have this question too, which is that your numbers sound low to them, so reconcile what Bick, Blandin, and Deming found versus what you found starting with telling us what they actually found in the first place.

NATHAN: So they ran sort of a current population style person-level survey, and they find that about 40 percent of the working age population reported using AI. And about 23 percent of the employed individuals used it in the past week or two at work.

And so those are obviously much higher than the firm-level use rates. And I think that this points to an interesting problem for the measurement agenda and thinking about how do we measure the diffusion of these technologies. And it goes back to this sort of incidental — it could be the case that these are incidental uses.

And so, workers would report, "Yeah, I used AI because I was bouncing my text off of it for typos or whatever else." But when the firm is thinking about, "Is the technology being used in the production of goods and services?" They might say "No." And so I think to reconcile these two things, we need to think about what it is that these individuals are using the technology for, and how does that interact with the firm's production technology.

CARDIFF: Here's another paper that you said was very cool and very interesting, and it was titled, *Answering the Call of Automation* from James Feigenbaum and Daniel Gross. It introduces the concept of "so-so automation." So tell us first what that is. What is so-so automation and why it matters?

NATHAN: I love this paper. The so-so automation, they're actually referencing some work by Daron Acemoglu and Pascual Restrepo who talk about so-so automation technologies as one that reduce the demand for labor because they're displacing workers, but they don't create a big enough productivity shock to create compensating labor demand in other places in the economy.

CARDIFF: So you only lose workers in the economy and you don't gain them in some other job?

NATHAN: Right. So if the technology displaces some workers but allows you to scale up significantly because of the productivity gains, then you might actually, on net, end up with more workers, right? Your labor demand on net might increase.

But a so-so automation technology is one where you're displacing workers, but there's not really a productivity effect. You're sort of taking those tasks and doing them, not any better, not any worse — maybe a little worse, but it doesn't actually affect your scale or your ability to scale or your labor demand.

CARDIFF: And so this paper answering the *Answering the Call of Automation* from James Feigenbaum and Daniel Gross, it looked at one of those so-so technologies, what was it and what was their finding?

NATHAN: That's right. So, they were looking at an automation shock of the introduction of automated telephone switches. And so in the 1930s, one of the most common occupations for women in the U.S. was a telephone switch operator.

CARDIFF: You see this in the movie sometimes—

NATHAN: That's right (CHUCKLES).

CARDIFF: You call and somebody at a switchboard—

NATHAN: Click, click, click.

CARDIFF: Picks up the phone, says, "Yeah, I'll connect you to James Street 554," or whatever.

NATHAN: That's right. And so AT&T in the 1930s introduced new technologies that automated that process. And, in short, what they found was that the incumbent workers, the older women that were working in those occupations, they did suffer losses in wages, decreased wages, less likely to be working in the future.

But the younger cohorts of women that came after them, they weren't permanently scarred by these technologies, which suggests that local labor markets were able to reequilibrate and adjust to the technologies. And firms were able to find new jobs and new ways to engage with those young workers without permanently reducing wages.

It's a really fascinating paper and it introduces this question of whether or not the differential effect on incumbent workers, and the process of how technologies are ingested by markets and how they adjust to it, and how those adjustments can be

quite costly for some people and some subsets of workers. But it doesn't necessarily have to be costly for everyone.

CARDIFF: Or costly in net terms for the economy.

NATHAN: That's right.

CARDIFF: So some people lose out. In this case, those incumbent older workers who lost out on those jobs of being a switchboard operator and overall the economy still adapted, even though you didn't get that huge productivity burst from this technology.

NATHAN: That's right.

CARDIFF: And so it leaves you with the question of, "Okay, so when there is a period of adjustment like that, do you do anything?"

And we'll talk about policy later, but I wanted to make sure that that was the setup for this paper, which is very interesting. Okay, next up is actually a website, AI-2027.com. It's by a series of people: Daniel Kokotajlo, if I'm pronouncing that right, Scott Alexander, and three other folks.

They basically go through some of the more, like boomy, not doomy, but boomy predictions of AI and how it's gonna accelerate economic growth. Is that right?

NATHAN: Yeah, there's this really interesting dichotomy in like the Silicon Valley AI types. There's the boomers and the doomers. The doomers are talking about AI safety and whether or not we're gonna be turned into paperclips and then there's sort of the boomers that think we're gonna get AGI, or an artificial general intelligence, or even artificial super intelligence, ASI, within the next couple years.

And so this website—

CARDIFF: What does that mean? Because it sounds like we're just giving adjectives here that add to the awesomeness of AI.

NATHAN: Yeah, no, that's what it is.

CARDIFF: (LAUGHS) Artificial super duper intelligence comes after, right?

NATHAN: That's right. The graphs, the intelligence graphs that they write down, it's not all that much. Once these systems get to a certain level of intelligence, it's sort of interacting with a God is kind of how they're thinking about it.

CARDIFF: So the essentially almost quasi-infinite economic growth is what you're saying?

NATHAN: Well yeah. So there's a lot of folks in this world that think, in this sort of milieu that think, "Oh, well we might get 10, 30 percent economic growth per year," and the group that wrote this put down this website. To their credit, they very nicely map out what that would look like at each stage, month by month, between here and there.

And I think, for us in the economics community, thinking about innovation and technology change from the economic modeling perspective, it seems kind of wacky and out there. But I think it's a really nice thing to keep on your radar as smart people are thinking about it in this way.

And it's something that I think it's worth engaging with.

CARDIFF: I was gonna say also for our listeners, for context here, economic growth for the last quarter century has averaged about 2 percent a year, real GDP growth, in the U.S. So going to 10 percent or 30 percent is just a massive massive increase — economic growth on an order that we have never, not just not seen, but not imagined previously.

So it'd be amazing. I think it is worth, like you said, keeping possibilities like this in your back pocket somewhere.

NATHAN: Yeah, that's right.

CARDIFF: Even if it's super low probability, like, "Hey, low probability things happen all the time," so—

NATHAN: In these uncertain times, it's worth keeping a wide lens.

CARDIFF: Right.

And then up next is a conversation between the podcaster Dwarkesh Patel and our friend again, Tyler Cowen, in which Tyler essentially is the counterpoint to those guys who are saying that the economy's gonna boom unlike we've ever seen. Tyler's

saying that actually, there are big obstacles between our world as it exists now and that world.

NATHAN: I think this was a really fun conversation that they had. And Dwarkesh, to his credit, he sort of embodies the uncertainty here. And he even admits on his show that he flip flops from week to week whether he's a boomer or a doomer. And so, he was trying to bring that argument forward and say, "Okay, why won't we get these extra normal growth rates, 30 percent growth rates?"

And part of what he was getting at is that these economic models that we write down in macroeconomics, they often turn on population growth. So, the change in the GDP growth is tied to population growth in some of these models. And so with AI, if you can sort of copy them millions of times, you get an effective population growth that's quite large and so shouldn't that flow through to economic growth.

And so Tyler, I think rightly, points to the fact that there's a lot of frictions in the system. There's a lot of bottlenecks. There's a lot of things both social, cultural, regulations, what have you, that are gonna slow down the diffusion of those technologies.

And that's consistent with the history of technology change, too, right? So when you got electrification of manufacturing plants, the whole plant had to be reordered. It's not gonna be an overnight process to implement.

CARDIFF: Stick with that example though, because that's a great example and it's one of my favorites, which is that in the late 1800s, with the advent of electricity, everybody's like, "Oh, this is gonna change everything right away!" And at first it did change some things meaningfully, so if you used electricity to change one part of what a factory does, you could see the improvement.

It was a nice change on the margin, but what was actually needed was a totally fundamental rethink of what a factory does in the first place and how it is designed, and that actually took decades. So you had this incredible technology and the potential to radically accelerate economic growth, but first, you had to get the human element right, which was to get people to try the technology in this radically transformational way instead of just using it to improve the things that we already knew how to do.

NATHAN: That's right.

There's a lot of experimentation that needs to happen. People need to think about how these technologies can be incorporated into their production processes. That's

not an easy task. And with AI, with how fast it's moving, that makes it even all the more difficult to sort of think about how to integrate these effectively.

And that would put us more into that world of "It's gonna take some time to diffuse," but there's also lots of regulatory and other barriers that you'll run into.

CARDIFF: And that brings us to our next example of something that people should read, which is an article from Timothy Lee who writes *Understanding AI*, which is one of my favorite Substack newsletters. Everybody should definitely subscribe to it. He wrote a piece titled "Driverless trucks are coming and unions aren't happy about it."

And I know you had a really positive reaction to this. You tell us the story because you love this story.

NATHAN: I love this story. It sort of embodies all of the most interesting bits of the bottleneck story. So, in the case of driverless trucks, the companies that were creating these new technologies were running up against regulations about the reflective triangles that could put out when a truck is disabled on the side of the road.

CARDIFF: Which, by the way, are required by law. In other words, if you drive a big truck, and trucking is a big occupation in the U.S., if you drive a big truck and you pull over to the side of the road, you have to walk out of the truck and you have to put down one of these reflexive triangles on the street so that people are warned that there's a big truck that has stopped and pulled over to the side. Okay, that's the setup. Keep going.

NATHAN: That's right. And so the companies that were producing these driverless trucks asked for a waiver on these regulations and they offered an alternative, which was a bunch of lights on the side of the truck or whatever, and they tried to prove their case that these lights were as good or better.

One of the funny little bits of that is, "Well, the triangles might get blown over by the wind and the lights won't."

CARDIFF: (CHUCKLES).

NATHAN: But the bottom line is that they did not get these waivers, right? And so now what are they doing?

They're thinking, "Well, how do we automate the process of getting these little triangles out? Maybe we have a little robot that comes out from under the truck. Or

we could pay somebody to drive behind the truck," you know? And then they'll be the triangle putter. This is one of those things where it's actually better understood as a proxy fight, right?

The people that are against the implementation of these technologies, they'll find these little places where you can create a bottleneck and accentuate a bottleneck.

CARDIFF: They use existing regulations to help them.

NATHAN: That's right.

CARDIFF: They'll use the law to help them find or introduce those bottlenecks so that it slows down the automation. And you see this in a lot of examples. There was the big fight over automation at the ports and the extent to which the unions are just really aggressively against automation there because they're worried about workers losing their jobs.

But the consequence is that there's no productivity growth at those ports, you know? And so, there's bottlenecks everywhere, even though the technology exists to make it all work better.

NATHAN: Yeah. And again, going back to the Feigenbaum and Gross paper. They're right to think that these automation technologies could hurt the incumbent workers, right? So that is gonna be the case. But that doesn't necessarily mean that we should be freezing the economy in amber.

CARDIFF: That brings us to our next example of something that you think our listeners should check out, which is a paper called "Generative AI at Work." It's by Erik Brynjolfsson and a couple of other scholars as well. Tell us about this one.

NATHAN: So this is a really neat one, and I think we need more of it. It's a paper that's looking at a pretty relatively narrow example of AI use and implementation. But you can start to learn some lessons that may generalize. And so they looked at the staggered rollout of artificial and generative AI among customer support agents, about 5,000 customer support agents.

What they found was that it increased worker productivity, but really interestingly, the least experienced workers got the biggest boost, right? So it actually decreased inequality among workers, so the return to those workers that had the least tenure of having this generative AI help them in prompting them to do their work, they saw the biggest boost.

CARDIFF: I wanna close with a discussion of what, if anything, we should do about AI, right? One of the topics that you like to delve into from time to time is the issue of "Should we try to steer this technology or not steer the technology?" So give us a flavor of what you mean by that and where you come down.

NATHAN: I think this is one of those cases where you can — it's easy to understand why people are worried about the displacement effects of these technologies. They could be quite large, but I think there's a couple reasons why it's not necessarily the best idea to go in and try to fine tune the system of innovation on the front end.

One is that these technologies are really nascent. There's a lot of uncertainty about how they'll change over time, and the history of technological change is — there's plenty of examples of technologies that were unexpected or the implementation of technologies went in directions that you wouldn't have predicted upfront.

And so when we start trying to put our thumb on the scale of the upstream innovation process, we might end up cutting something off that would've — we might say, "Oh look, this thing looks like a so-so automation technology." And so we either try to regulate it or stop it in some other way.

But that thing might actually, if you let it bloom through the system, might actually generate some labor demand boosting technology later on. There's also parts of the economy that I think we would love to have even so-so automation technologies. So, there's this wonderful idea in economics that some of my GMU folks have written about is Baumol's Cost Disease.

And so it's the idea of something like when there's a sector that has seen really slow productivity growth relative to other sectors, the labor costs in the low productivity growth sectors will go up, even if it's that sector has low productivity growth because their outside options are better.

Right? So they could—

CARDIFF: Of the workers, right?

NATHAN: Yeah, the workers in that low productivity industry, they could go work in the high productivity industry—

CARDIFF: Which can afford to pay the high.

NATHAN: —and get the high wages, right?

CARDIFF: So the only way that a low productivity sector can afford to pay those workers to join is if they raise the price of their product. So inflation in those sectors.

NATHAN: That's right. And so, in those cases you might actually say, “Well, no, I'll take the so-so stuff. I'll take those technologies that displace labor without boosting productivity because it puts much needed downward pressure on prices.”

CARDIFF: Very interesting. The examples that you mentioned where we didn't expect the technology to go a certain way, but it did and it was amazing. Tell us one or two of those.

NATHAN: So there's this wonderful one with the laser. So the gentleman who invented the laser famously said, “It’s a problem in search of a solution.” Now, it's embedded in lots of communication technologies and many other things. The inventors of the early versions of touchscreens, right? So they were, I think, originally invented for air traffic control systems. And what did we get? We got Angry Birds, right? So there’s lots of ways that innovation unfolds that are really surprising.

CARDIFF: To be clear, you're being complimentary about Angry Birds, right?

NATHAN: Yeah. Well, more or less.

CARDIFF: More or less, kinda. (LAUGHS) not about Angry Birds specifically, but about the other uses of that technology.

NATHAN: Yeah. The point is these things are really unpredictable and the way that innovation unfolds is unpredictable. Brian Arthur has this nice description of innovation as like an emergent process, right? It’s combinatorial, like the things build upon one another.

One innovation that comes in, builds on others, and you combine the new things that you have in new and complex ways. It's complex, it's combinatorial, and it's emergent. And so for all those reasons, it's really hard to predict which direction it's gonna go.

CARDIFF: This happens in medicine a lot too, where you start taking medicine for one thing and you don't realize it has like these other effects. In science, this happens quite often. Wasn’t penicillin sort of an accidental discovery too?

NATHAN: That's right. A Petri dish, right? Sitting on the windowsill or something like that?

CARDIFF: Yeah, something like that. It's fascinating and I think, if I'm understanding your point correctly, you're saying that if we overregulate AI at its onset, when it's really starting to become popular, that there's a risk that we end up missing out on the fascinating new directions that it could go in and all the gains from that.

NATHAN: That's right. And it's also the case that, I think more to the point is that we just don't know. We don't know how the general equilibrium effects are gonna play out. There's lots of predictions about the labor market effects and the displacement of workers.

We don't know how firms are gonna react and we don't know how labor markets are gonna adjust, and we don't know how the innovation ecosystem itself is gonna change and adapt to the changes from these new technologies. So it's just very, very hard to know what direction it's gonna go.

So if you start putting pressure on the upstream innovation system, you might be cutting off things that would really benefit humanity.

CARDIFF: And a lot of times, I'll read a lot of things about how societally or policymakers should respond to the rise of AI and precisely because of that uncertainty that you just mentioned, it seems like we always end up with a policy that you can have a debate about whether it's good or bad, completely, regardless of whether or not it's a response to AI.

So you mentioned that example earlier of the workers who lost their jobs as switchboard operators, where the economy on net still benefited from the introduction of this new technology, but there was a subset of workers who did suffer, and so maybe you put in place the right safety net to help those workers.

But that's not really a conversation about technology anymore or about AI. You could make that argument about any worker who loses their job for any reason or is suffering for whatever reason. You could have these debates about how to organize society and how to put in place a good political system that's responsive to people being out of work or suffering particular labor market effects completely separate from the conversation about how to respond to AI, which maybe that conversation should exist in the realm of cybersecurity or something like that, as opposed to a labor market response policy.

NATHAN: Given all this uncertainty, I think you wanna focus on policies that are gonna be a good idea either way, right? So there's maybe some physical infrastructure change things that you could be doing, you could be hardening

infrastructure. You could be in increasing energy transmission, all these other physical infrastructure things.

Maybe those things complement the use of AI, but they might be good ideas in and of themselves. And you could say the same thing of the labor market impacts. So, you might come up with policies that you might want to have, whether or not AI turns out to be labor displacing or not, right?

So you wanna focus on those things that are gonna be a good idea either way. Why? Because we have no idea how it's gonna shake out.

CARDIFF: Nathan Goldschlag. Did you have fun on your first post-government podcast?

NATHAN: Yeah. This is a little different than what I used to do in the data mines of Census, but this is a lot of fun.

CARDIFF: Good. Well, it was great to have you. We should go back to work now.

NATHAN: Yeah. (LAUGHS)