

Goldman Sachs Exchanges
How Will AI Impact the Labor Market?
Daron Acemoglu, Institute Professor, MIT
Neil Thompson, Director, FutureTech Research Project,
MIT
Joseph Briggs, Lead, Global Economics Team, Goldman
Sachs Research
Allison Nathan, Senior Strategist, Goldman Sachs
Research
Date of recording: June 2026

Allison Nathan: Rapid improvements in AI capabilities and growing corporate adoption have led some prominent technologists to predict that AI could eliminate a massive number of jobs before the end of the decade. So, just how concerned should we be about an AI job apocalypse?

I'm Allison Nathan, and this is Goldman Sachs Exchanges.

Each month, I speak with investors, policymakers, and academics about the most pressing, market-moving issues for our Top of Mind report from Goldman Sachs Research.

This month I spoke with MIT's Daron Acemoglu and Neil Thompson, as well as with Joseph Briggs, who leads the global economics team in Goldman Sachs Research.

I started by asking Joseph just how much labor displacement from AI he expects ahead.

Joseph, I know you are well aware that there's been a lot of debate about AI's potential impact on the US labor market, especially as we've seen some companies citing AI as a factor in recent layoffs. We've all seen the headlines. What are you expecting in terms of AI-related labor displacement in the near term and over the longer term?

Joseph Briggs: So to level set, if we look at the labor market today, you can see the imprint of AI in a few industries and a few sectors where we know that AI is already having an impact. And so if we combine across sectors like tech and management consulting and graphic design, areas where tools have already been developed and deployed, you can see that overall, there's probably a 10 to 15,000 drag on month-over-month job growth from AI impacts. All that being said, it's still a fairly narrow labor market shock, and we're not seeing a big impact today on the broader economy.

Now, I do expect that will change going forward. Under our baseline forecast for a 15% uplift to productivity following full adoption of AI, if we combine that estimate with the historical elasticity between how much does a technology-driven productivity increase tend to displace workers, we come up with an estimate that around nine percent of all workers in the US will be reallocated to new positions during the AI transition.

Allison Nathan: That's a pretty big number.

Joseph Briggs: It is a big number. Nine percent of workers being displaced by AI would correspond to 15 million workers leaving or being displaced from their positions today and having to find new jobs. Displacing nine percent

of workers would be the type of automation and reallocation shock that we saw in the late '90s and early 2000s and in other periods of significant technological change.

What I'd really emphasize is that it is over a 10-year period and as long as the displacement and the job loss is spread out enough, then the impact on the unemployment rate in any given year will likely not be that large. So for example, even under our forecast for nine percent of workers being displaced, we'd still expect that the unemployment rate increase in any given year would be less than one percentage point.

Allison Nathan: And then also a key part of your forecast, is that in addition to these job losses, you'll have creation of jobs. So talk us through your assumptions and expectations there.

Joseph Briggs: Yeah. We're not expecting that displaced workers will be displaced over the long run. We do expect that ultimately there are going to be more than enough jobs created to reabsorb workers back into the labor market. And the key reason for this is that there's a long historical record of technology delivering significant job gains.

A couple stats that I would flag around that, if we look back over the last 80 years, around 85% of job growth has been driven by the technological creation of new positions. Likewise, the U.S. labor market is incredibly dynamic.

Every year, we see around 30 million jobs being created, now granted, 29 million are being destroyed, all the time, implying that technology and automation are constantly

leading to a significant amount of labor market churn, where new positions are created and jobs are destroyed.

Now, we think that this will repeat itself going forward. Particularly in a world where AI is enabling innovation, even a five percent acceleration in the pace at which new jobs are being created, would be more than enough to reabsorb the workers that we're expecting will be displaced by AI-driven automation.

And so, you know, over the long run, I'm really not concerned that we're going to see permanent job losses. The bigger question is: does that pace of new job creation pick up fast enough to offset any term headwinds?

Allison Nathan: Just to be perfectly clear, you don't subscribe to the view that we are going to see a world in which a lot of people just don't end up with a job?

Joseph Briggs: Yeah, definitely not. I think that the view that has been put forth by a lot of tech commentators, where jobs are going to be permanently displaced, it really focuses on the job loss aspect, which as we've discussed, will likely be pretty meaningful, but it ignores the job creation aspect. And as long as we see history repeat itself and we see that technology does, again, as it always has, lead to new work opportunities, then we won't see permanent job loss over the long run.

Allison Nathan: MIT's Neil Thompson is less convinced that AI will displace a large number of workers. He argues that capability alone isn't sufficient to do so, and he points out that jobs consist of many tasks, only some of which can be automated. So, he seems to expect a slower and

more uneven labor market adjustment than AI's current capabilities suggest. And he also takes some comfort in the idea that we can see AI coming. Here's some of my recent conversation with him.

Neil Thompson: We should absolutely think of AI as being this very transformative technology that is not only very capable, but actually becoming capable very quickly. But when we then want to connect that to jobs, it's really important to say that AI capabilities are only one in a series of steps that lead to a change in jobs. Right. And so you first say, okay, could AI do this task if it was given all the right information. Really crucial in there was that if you give them all the right information, right? If you actually think of lots of things you might imagine doing, so you say, oh, well, maybe when you check in for a doctor's office, some of this could be done. But of course, as soon as you want to get any kind of medical advice, all of a sudden you have to get access to privacy records and things like that. And so you can very quickly get in a situation where, oh, you need these kind of records. So, there's a whole bunch of stuff needs to be done there.

Then even if you say, ok, now I know how to build such a system so that it can get all the right information. You can ask the question, is it cost effective to do? And some of the previous work my lab has done has shown that in many cases it might not be because you might need such an exacting system that it would cost a lot to run. And so, you really need all of those pieces to come into play. You need

AI to have the capabilities. You need to be able to provide it all of the information it needs to make those decisions. That's not an easy thing in many cases. And then you need to know that once you do all those things, it will be economically attractive in order to have the effect.

So what that in practice means for lots of people thinking about this process is that we're going to look at capabilities and they're going to improve very fast and we're going to say, my goodness, AI can do a lot. But then there's going to be this adoption process that is going to take a much longer time. And that means that large businesses are going to get automated before small businesses. It means that there are things that are more important and more attractive to be using AI for are going to get done before a very long tail of things that probably will take a long time or not happen at all.

And this actually is not that different than what we've seen in previous waves. The previous waves of automation that we had in say the '80s looked at automation as what can you do that is routine tasks that you could imagine building into a pipeline of, say, what a computer can do. And there were quite a number of tasks that we said, oh, we could imagine doing those but in fact, only a fraction of those have been automated. And so, the question for us, and we don't have an easy answer to this, but it's how fast that adoption pattern will happen. But it certainly is going to take a lot longer than the growth of AI capabilities.

Allison Nathan: Neil, you often talk about this in terms of expert versus inexperienced tasks. Why is that distinction important here?

Neil Thompson: So, this is very important because in some of my collaborations with David Autor, one of the things we see is that most jobs, it is not that all of the tasks in that job are going to be automated. In most cases, we're going to see partial automation of jobs. And then the question is, if your job is partially automated, what happens to you? And our intuition on this is often on the demand side, by which I mean that if we think of, if 30% of my job got automated, that might decrease the demand for my job by 30%. And when we think about that, we think of prices should go down and quantity should go down. So that means I should get paid less and there should be fewer people doing my job.

I think a better intuition for thinking about this is on the supply side of things, which says, if somebody automates part of my job, what happens to me really depends on what the task is that gets automated. So if you say filing my expense reports, right? That is a very inexperienced part of my job. I'm pretty happy for someone else to take that and focus more of my time on the stuff that really makes me valuable. So intuitively, we can say that sounds more attractive to me. Whereas if you have a system that comes in and does the most expert part of my job, and I'm left doing more of my expense reports and stuff like that, that doesn't sound like it's an attractive deal to me.

And so, we can actually analogize this to what has been happening over the last 30 or 40 years. For example, taxi drivers, so when GPS comes in, it automates the most expert part of what a taxi driver does, which is knowing all of the routes around the city. That means that what actually happens there is now all of a sudden many more people can do taxi driving. There's a lot more competition that drives wages down. So, wages do indeed go down. But in fact, there are many, many more taxi drivers now than there ever were, they would just call them Uber drivers.

Conversely, if you think about proofreading, proofreaders used to do spell checking before we had spell check, right? And that was not very expert part of their job. That, of course, has been completely taken away by Microsoft Word and all of those others doing the spell checking for you. But what's left in the proofreading job is the much more complicated, like, how do you think about structuring an argument? How do you marshal evidence in the right way? So lots of people could check spelling. Not that many people are really good at the other part. And so as part of things got automated, there were fewer people who could do it, but the people who did do it actually got paid more. And so you can see this pushes us in two different directions. If your least expert stuff gets automated, you become more expert, your wages go up, but there are fewer of you. If your most expert stuff gets automated, that actually pushes your wages down, but actually more people

enter. And that means that there are typically more jobs, not less jobs.

I think that we will face the same thing with AI. It will, in some case, automate the more expert, in some case, automate the less expert. And so that's going to mean that there's going to be a diverse set of labor implications for AI. And this is actually a very, very important thing that governments and businesses need to think about because it means that if they're just planning for a uniform, everybody has a bad outcome, right, it's actually a much more nuanced thing.

Allison Nathan: But Neil, let me just ask the bigger question, which is what does all of this mean for how AI will ultimately impact the aggregate number of jobs?

Neil Thompson: So I think it's an important question, but it's a question that if we think about it at that level, you have a couple of things. So you have this expertise effect as you have a partial automation of a job. And that, as I say, does not push us particularly in one direction or the other. But at the very aggregate level, what really matters is there will be presumably some jobs that will be largely automated. And so those will actually disappear in some sense. But of course, new jobs and new tasks are also going to be created. And that balance between the jobs going away and new jobs being created, we really don't know. If we look at previous automation waves, we see that in general, there are lots of new work that comes about,

right? So, in general, this has been okay. Now we know that the way that AI processes things is more similar to humans and therefore we might be a little bit more worried about that.

But I think it is really too early to know whether we should be expecting a lot of unemployment from this or whether the new tasks we have created and the extra leverage that humans will get will actually be very important in allowing us to still have lots of work out there.

Allison Nathan: But AI's capabilities are much broader than many past technologies. So, does that make significant job replacement more likely than in the past?

Neil Thompson: So I do think that AI as a technology is considerably broader than many technologies we've had in the past. That also comes with some aspects of it that are more difficult to implement. The way I think about this is if you think about a traditional technology like, databases or Excel or something like that, it has a pretty limited scope. But within that scope, it performs basically at 100% effectiveness? You never worry about Excel multiplying numbers wrong or your database missing a third of the records or something like that. If you put in the right query, it's going to give you the right answer. So it's narrow, but very effective. AI is much, much broader. You can ask it a multiplication question or you can ask what to have for dinner tonight. It will give you answers to both of those questions. But its ability to get something 100%

correct is much more limited, right? It's much harder to stop there from ever being a case of hallucination, ever going off the rails. That means that it's very apt to be used as a tool that can help people, but it's harder to use it as a tool where you can have a modular part of your process that you can just forget about. So yes, it's broader, but there are also some other aspects that are challenging for it. And that same thing that makes it very broad also makes it much more powerful for being a tool to do stuff. And of course, it's important to remember this happens at several different levels? So it happens at the level of, I'm a worker, someone gives me this tool, I now do my job a little bit faster. But even if you say like, okay, in my organization, maybe there used to be five people. And two of the roles disappear, but that augments the rest of us. And that means that we now hire more people and business gets bigger faster or something like that. And so there are lots and lots of different effects here. And so I think we should be very skeptical about saying that it's going to destroy work and not create work along the way.

Allison Nathan: So, are the worries that AI will lead to a job apocalypse, as we've been hearing, warranted, or overblown?

Neil Thompson: So I think that people are right to look at the AI capabilities evolving and to say this does present a potential challenge to labor. But in one of our recent papers, we talked about the difference between crashing waves and rising tides as to how it affects human workers.

And I think this is important because if we think about crashing waves, you can think of this as like everybody in the workforce is walking along the seashore. And we're all like, ah, it's a beautiful day. We're not wet. We're just warm and sunny. And then a wave comes out of nowhere and a bunch of people just get totally washed away. That's a world where it's pretty anxiety producing for workers. What we see in our research is that does not seem to be the dominant way that comes in. The dominant effect seems to be rising tides by which we say like, okay, we're still all the seashore and some of us are on the sand, some of us are up to the ankles and some of us are up to the knees. But as the tide comes in, you say, okay, now it's a little deeper, it's a little deeper. So like AI is coming but it's not totally unexpected. And so at least if we're paying attention, we can have a good sense of what AI is going to do. That doesn't protect us from a fast rising tide, but it does mean that we can see it coming and it won't be as big of a shock and we can manage that process better. And so I think that to me is an encouraging sign of businesses and workers can look at what AI can do, can try and manage that process in a way that is much more active than if it was a crashing wave scenario.

Allison Nathan: MIT's Daron Acemoglu sees it a bit differently. He expects AI to have a small net negative impact on labor over the next several years. But he warns that job losses could be larger over the longer term if AI investment continues to focus more on replacing workers

than on complementing them. Here's some of my recent conversation with him.

Daron, let me first ask. The consensus among economists seems to be that the aggregate labor market data isn't indicating a significant labor market impact from AI yet. So, do you expect a larger impact to be visible in the near future?

Daron Acemoglu: It's very, very difficult to make any kind of predictions with any degree of certainty. But I would imagine that in 2027 we would see a little bit more of layoffs or slowdown in hiring in jobs where AI can at least have a chance of replacing some tasks. I do not think that we are currently seeing models and capabilities that are that good at complementing workers yet. I think a lot of workers are using AI for small things like checking text, some reference checks. And that's fine, but that's not like the really big complementary uses of AI. So I wouldn't expect that except in a few fields like, say, biological research or chemical research, I don't think we're going to see huge uses of AI in a complementary way that quickly. So I would expect some net job losses, limited, but net job losses within the next five years.

But, I want to also underscore that none of this, in my estimation, will be of a scale anywhere close to the kinds of things that some expect. I would say less than two to four percent.

What makes me cautious on the spread of AI is that we don't have easy to use applications based on the foundation models that can be adopted by many large-scale employers. I think you would need more reliable applications built on the foundation layer for these jobs rather than individuals or managers themselves prompt engineering, which would be inconsistent, time consuming, often challenge the knowledge of teams to use AI, unreliable, all of these things would make that not as likely a scenario. So, then we would really need to rely on these applications and we're not we are not seeing those applications yet. Coding, software engineering is an exception because, first, the models are already quite capable in coding because parts of coding have now been quite routinized. And second, the people in the software engineering space are all quite experts in AI, so they can give the right prompts and then troubleshoot and check the work. So that's why I think coding may be an exception in that some impact can be seen without these kinds of reliable applications, easy to use applications. But again, there's a lot of uncertainty.

Agentic AI opens the way to develop more of these applications. But once again, I don't know that it's that likely or that productive if every company has to develop those applications themselves. So the agentic advances would be most useful if AI model developers or AI application developers could use those to offer to the market reliable, flexible, easy to use, tools that other

companies can then adopt. So like the Microsoft Office version of AI, so to speak.

Allison Nathan: What type of worker is most vulnerable?

Daron Acemoglu: At the moment I see still the most vulnerable tasks to be those that are cognitive and routine, meaning that involve similar things and not too many new things, not too much innovation, not too much creation and not too much intense social interaction or judgment. So those would be tasks like customer service reps or back-office work. There are a lot of workers who do that. I think in total, if you add those two, it will be eight million in the United States, nine million. So that's not a small amount but also not huge.

Allison Nathan: But does that estimate of AI's labor market impact shift over the medium-to-longer term?

Daron Acemoglu: It's much, much more uncertain over let's say 10 to 15 years, it will depend on where the investments go. I have argued for more than 10 years now that the complementary path is actually quite productive. It's just that we haven't invested in it. It has a lot of preconditions. It has a lot of different investments that it requires, either at the pre-training level or the application level. It requires very different kind of data, very high quality data, but much more domain specific. I believe we are not making those investments sufficiently. So, I would

expect bigger net job losses within the next 10 to 15 years if things continue like this.

There are also so many big wildcards. The big, big, big wildcard, is the integration of AI and robotics. There are efforts on this much more outside of the US than the US, but also in the United States. And if there were amazing breakthroughs there, that would open up the number of jobs that AI would impact hugely. Physical tasks is about 50% of the work in the US economy.

The next layer is jobs that involve judgment, middle managers. That's where the agentic advances are going to be very important and the applications are going to be very important. That is very uncertain but we may get more information on that in the next year or so. And then the next big chunk is jobs that involve social interactions, where you need two kinds of changes for those to be in the crosshairs. One is the models need to improve in the social dimension. They're already not that bad in some of those, but they would still need quite significant improvements. But also that human consumers need to change their preferences and what they see as normal, that they accept social interaction from AI bots more than they do. And the new generation is more open, I think, on some of these issues. But how quickly that will go, I don't know.

Allison Nathan: But let me zoom out for a second. The vast majority of job growth since 1940 has been driven by

technology as well as the creative destruction process. So, do you think this time will be different?

Daron Acemoglu: There is no general law of economics that says that job creation always has to match job destruction. If you look at the last 80 years, a lot of job growth has come from changing tasks, new tasks, changing structure of occupations. But it has not been at an even pace. We have not generated as many jobs for workers without a college degree since 1980 as we used to before 1980. And that's why if you look at the employment to population ratios of especially men without a college degree, they have fallen quite a lot. And wages have fallen and stagnated for workers of that sort. And during certain periods like the 50s, 60s, early 70s, creation exceeded destruction and created a lot of demand. So that's why wages increased even faster than productivity during that period. But it has fallen short of destruction since the late 1970s with some periods of exception in between. If we just were to repeat the period since the 1970s, but now with the destruction coming from cognitive office jobs, that would already be, I think, the kinds of limited job losses that I'm talking about. So, I'm not saying that there is anything completely different this time. But every episode is different because those balance between creation and destruction, automation versus new tasks and changing occupations, those are different in every sub period.

Allison Nathan: And so Daron, what do you think the impact of these shifts from AI will be on wages and income inequality, which is very much in focus right now?

Daron Acemoglu: First of all, I think that inequality and employment are actually more tightly linked than sometimes is implied. If you experience declining or stagnant wages for some groups, they typically also reduce their employment population ratio or the participation rate. That's why I believe that some groups are going to suffer in terms of their wages as well as seeing somewhat slower employment growth or some employment declines. That therefore is the basis of my belief that we should expect increase in labor income inequality.

Now, there is one caveat to that, an important caveat, which is that if the jobs that are replaced were done by managers or already highly paid workers, that would work out differently, that now you're not replacing the jobs that blue collar workers used to perform like the technologies of the 1980s, 1990s, 2000s, which then increase inequality because those middle class type of wages were replaced by lower wages that many of these workers could get only in lower ranked occupations. If somehow you started replacing managers in large companies who are well paid, inequality could decline. I don't think that's the most likely scenario because my earlier account, it's the simpler, cognitively more predictable jobs customer service, back office, those are not the highly paid workers. And moreover, very well paid employees would go and find jobs in other

occupations. So they wouldn't be the ones that bear the burden as much as then the next layer who are then displaced. So that's the basis of my belief that labor income inequality would increase.

Allison Nathan: Honestly, I'm not sure if I came away from these conversations more concerned or more comforted. There's a lot that we still don't know. But I'll leave it there for now. My thanks to Daron Acemoglu, Neil Thompson, and Joseph Briggs.

And thank you for listening to this episode of Goldman Sachs Exchanges, which was recorded in June 2026. I'm Alison Nathan.

The opinions and views expressed herein are as of the date of publication, subject to change without notice, and may not necessarily reflect the institutional views of Goldman Sachs or its affiliates. The material provided is intended for informational purposes only, and does not constitute investment advice, a recommendation from any Goldman Sachs entity to take any particular action, or an offer or solicitation to purchase or sell any securities or financial products. This material may contain forward-looking statements. Past performance is not indicative of future results. Neither Goldman Sachs nor any of its affiliates make any representations or warranties, express or implied, as to the accuracy or completeness of the statements or information contained herein and disclaim any liability whatsoever for reliance on such information for

any purpose. Each name of a third-party organization mentioned is the property of the company to which it relates, is used here strictly for informational and identification purposes only and is not used to imply any ownership or license rights between any such company and Goldman Sachs.

A transcript is provided for convenience and may differ from the original video or audio content. Goldman Sachs is not responsible for any errors in the transcript. This material should not be copied, distributed, published, or reproduced in whole or in part or disclosed by any recipient to any other person without the express written consent of Goldman Sachs.

Disclosures applicable to research with respect to issuers, if any, mentioned herein are available through your

Goldman Sachs representative or at
<http://www.gs.com/research/hedge.html>

Goldman Sachs does not endorse any candidate or any political party.

Copyright 2026. All rights reserved.