Goldman Sachs AI Exchanges

CIO Marco Argenti on the future of AI in the workplace

Marco Argenti, Chief Information Officer, Goldman Sachs

George Lee, Co-Head, Goldman Sachs Global Institute

Allison Nathan, Senior Strategist, Goldman Sachs Research

Date of recording: March 6, 2025

Allison Nathan: Welcome to Goldman Sachs Exchanges. I'm Allison Nathan, and I'm here with George Lee, who is the co-head of the Goldman Sachs Global Institute. Together we're co-hosting a series of episodes exploring the rise of AI and everything it could mean for companies, investors, and economies.

George, good to see you again.

George Lee: Great to see you. I really enjoyed our first discussion and super excited about this one today.

Allison Nathan: So today we are going to discuss how AI can fit into enterprises. We are always talking about the

use case and perhaps we should say it's a big question about how enterprises can fit themselves around this evolving and exciting technology. So set the stage for us. When you think about what companies are doing today with AI, how would you explain the usage and how has that evolved over the last year or two?

George Lee: Yeah, great question. Well, look, I think we're in the early days and that's consistent with the state of development of the technology itself, which is quite novel as you and I have discussed before. That having been said, there is definitely a value gap opening up between the trajectory of the technology and the amount of enterprise adoption of that technology, and I think there are really three reasons for that. First is: Always was it thus? And by that I mean that the diffusion of new technology in businesses is always a little bit of a slow, deliberate process. Second, the pace of progress of the technology itself is an obstacle to adoption. Any technology that is moving as fast as this one makes it hard for CIOs, CDOs, etcetera, to figure out what technologies to deploy, when, where, and how. And so almost it's funny, as much as I admire their trajectory of improvement for the sake of adoption, I almost wish it would slow a little bit. Third, it's

a general purpose technology. That's part of the genius, but that's part of the challenge for an enterprise. So its use cases are bounded only by human creativity and imagination, but it takes human creativity and imagination to determine how you want to use this very broad technology.

Now, I have two caveats to that. One is there's like a generational component to this. In the department of, it's hard to teach old dogs new tricks, as new people join our organization and organizations like ours, they inherently bring with them an appetite to use the technology, a facility with the technology, and they will, in a way, drive adoption naturally. So that will happen. The other thing I would say is there's a big distinction today between big incumbent companies who are using this technology or attempting to use this technology and de novo enterprises. If you're starting a business from scratch today and you can write code unfettered by legacy, etcetera, you can instantiate processes that are built from first principles around AI.

It's an amazing piece of leverage. We have encoded workflows, habits, regulatory obligations, new technology, old technology, and fitting this new tool into that mosaic is

just inherently harder. But man, if you look at the data from, I saw something from Y Combinator recently, new businesses being formed today that are using this technology to build the software platforms, allow them to operate are just zooming, they're moving really fast. So, look, early days, I'm very optimistic, and it be fascinating to talk to our guest about the actual practical day-to-day challenges of deploying this new technology.

Allison Nathan: So let's bring Marco Argenti into this conversation. Marco is our chief information officer here at Goldman Sachs. Thanks so much, Marco, for joining us.

Marco Argenti: I'm happy to be here. And, this is definitely an exciting topic for us to discuss.

George Lee: Great, welcome

Marco Argenti: Thank you.

Allison Nathan: Marco, given the context that George just gave and the many developments in gen-AI over the last even weeks, certainly months, has your thinking about AI changed at all in this time period?

Marco Argenti: I think the world of AI has changed like three or four times over, even in the last year or so. And I

think what's interesting is that every time we seem to have reached some form of a plateau, here come sort of a bolt from the skies that just kind of makes us rethink any illusion of the technology having reached any sort of scaling. So, if you look at even like the last two months, there has been some of the most fascinating evolution with regards to specifically thinking models and models that have the ability to reason. And if you look at those are not incremental changes. Also for the user standpoint you now have models that, for example, can do real research for you in a depth that was unprecedented before. I personally tried it. I compared all these models on a topic that I know really well, which is kind of the last few years of the history of Nokia smartphones, because I used to work there and it's an interesting like obviously case study of strategic choices. And I had six different AI doing the research and some of that came out as the best book I've ever read on Nokia, it felt almost like it was written by an insider. And then we tried other subjects. So I think what I'm trying to say is I've been in technology like sadly for about I would say since I wrote my first line of code now 45 years, almost exactly. And, I've been through revolution such as the beginning of mobile or the internet, or the beginning of the

cloud, at Amazon. This is compressed in a way that I've absolutely never seen before.

Allison Nathan: So we all agree that the advances the technology itself is making is absolutely enormous and unprecedented in terms of the speed. But when you think about this topic of enterprise adoption, I mean, George already said it's not necessarily keeping pace, but what are you observing from your seat?

Marco Argenti: I think, in general, I agree with the fact that enterprises in general are lagging the state of the art with regards to adoption. Also, because, I mean, let's face it, it took, even like decades to be perfectly proficient with computers. And then it took, you know, years to be implementing the internet in a meaningful way. It took a couple of decades even for adopting the cloud. I would say we're like in year one-and-a-half of actual useful products. Right. And not toys. So let's give ourselves a break a little bit, because guess what, on one side you have a brain that is evolving on a daily basis. On the other side, you have a human brain that- and habits- that generally tends to evolve in hundreds of years or millennia, biologically and at least years from a habit perspective. And so I think right

now the biggest friction point to adoption of AI in the enterprise is people, is behaviors, is the fact that you have muscle groups that you need to retrain. Although I have to say that we're starting to see areas where there is actually fairly consistent and meaningful progress. And those are in areas that may be are a little bit easier to change because the people that work in there are more in a way, open to change. And one of them, for example, is the developer community.

The developer community in many cases tends to be eager to experiment new things. Historically, they know how to deal with imperfect products. A new operating system is released 3, 4, 5 versions to developers before it actually hits the market. And so you are kinda know how to use stuff that is not there yet and you kinda use the good stuff and then you know how to deal with the bad stuff. And so I think that community has been very receptive. And that's one of, you know, the first area that in Goldman has gained real traction of really improving our developer experience. And then as a result of that, gaining more efficient. There are many other cases, but I will say that today, the real answer to your question is this is probably the one of the

biggest change management challenges and exercises that any corporation I've ever seen in their history.

George Lee: Totally agree. And Marco, I'm wondering one of the things I talk about with this technology is, especially as it relates to adoption, there is this magic button hypothesis. The idea that our users tend to expect it to work perfectly and solve all their problems, and unless and until it does, they tend not to want to adopt.

Marco Argenti: Yeah.

George Lee: Again, caveat, as you write, certain populations more inclined to do it. How are we getting people on the journey with us and how do we identify people who are willing to experiment, to change habits, to reconsider. Talk a little bit about what we're doing to foster adoption by finding those people inside the firm.

Marco Argenti: You know, if you want to have a disruption or if you're kinda be ready for a disruption, you need to find people that have already crossed the bridge of accepting that a disruption is going take place and be very open to it. And so we have really started to seek for disruptors that are- have taken a very kind of purposeful approach on how to navigate through that disruption and

ending up on the winners side, even at the sacrifice of questioning their own beliefs or even sometimes really profoundly rethinking their own businesses. And so we identified quite a few of those. And the trick when you're finding those disruptors is to find people that are actually in a position where they can drive change. And I think we have a few you know, I don't want to, you know, call out names, but it's starting to become a movement and those people in turn will influence others who will become role models, and then they will actually apply their judgment. So because we don't want reckless disruptors, we want mindful disruptors. And I feel pretty good about how we've been doing so far.

George Lee: Exactly. Now, I think the dynamic inside businesses, once you identify those disruptors, they're willing to take some risk, I think from first principles, they start to generate positive results. There's a follower effect in in enterprises, and so that's what I think we're hoping to catalyze. Marco, you talked about one of the obvious use cases, which is instrumenting our developers with this technology. I can think of a, a few other sort of more pedestrian use cases. Let's dream big for a second. What are some of the most ambitious use cases of this

technology for knowledge intensive business like Goldman? What are, you know, let's fast forward a couple of years and what are some of the inspiring use cases you can imagine in a business like ours?

Marco Argenti: I think my belief is we're heading towards a world where we're going to live in a hybrid workforce where you're going to have to have the ability to manage interactions with human colleagues and AI colleagues with the same ease. And so in this hybrid workforce, one of the great advantages is the fact that you can achieve elasticity, the same way as you do with the cloud. If 10 years ago someone would ask, asked the CIO of any large company, how many servers you have, they would've said the number with a very great precision. Today, if you ask me how many servers I have, I tell you I have no idea because it might change on a minute by minute basis because with the cloud, if you have a high workload, especially in the world of serverless, you can go up and down in minutes, and I think that is actually going to be the same answer if you ask me 10 years from now, because you could actually surge agents for a closing of a quarter or for an earning season or for particular times in the market, and then shrink, and then go up and down in the agent side where

you know your human side you can remains more stable. And talking about, you know, the way this is going to impact our daily life, I think, one of the first objectives that we kind of already, I think, kind of realized is that pretty much every developer right now has AI in their what's called id, which is their interface to do coding. Okay. I want that to happen to every single employee, but not for coding, but it should be basically as natural as Google. Okay? It should be as natural as doing a web search. It should be actually most importantly, as natural as sending an email to someone asking, can you do this for me? Or can you find someone that could give this information, like client report or research, etcetera. So going into really a world where you can ask natural language questions and you get very complex, very GS-specific answers. And for that, we built and we're deploying at scale our GS AI assistant, which is really like your desktop for AI. That's the way to think about that. And from there is where you go to find answers. And I think, to me, the evolution and what I see as a really fascinating future is where this assistant is actually becoming and sounding more and more like a very experienced GS employee or a very deep expert on a certain subject, rather than giving you generic answers. Because at

the end of the day, the companies that are going to win are the companies that are going to have a workforce that is not only embracing AI, but also, in a way knowing how to make AI part of their daily life and their daily work.

Allison Nathan: Marco, there are some guardrails that an institution like Goldman Sachs, you know, has to be aware of and, put in place. So, I mean, talk to us about what we might be doing to try to overcome the limitations that we face in the type of information we deal with.

Marco Argenti: We are using this, the latest models, but we're very careful at controlling you know, any potential infiltration of external data or data externally. And most importantly, we are doing a lot of work to, you know, quote unquote ground those AI so that they don't make stuff up. That is the most important thing.

Allison Nathan: Right.

Marco Argenti: So rather than removing stuff, we tend to actually make sure that stuff is actually accurate by cross-checking it with sources.

So I think, you know, that's one of the things that, unfortunately we live with when we talk, about AI that, you

know, when you use AI, as you know, the wild type of AI that is out there, it's kind of a medical term, which means what's normally out there. You in many cases tend to get information that is not necessarily accurate, but it sounds really plausible. And then if you kind of try in a sort of a heavily grounded AI, it might sound a little bit less creative or imaginative, or maybe it's more correct. And so I think we've been putting a lot of, our biggest effort there has been to increase accuracy, reduce hallucinations, and also to make sure that, we minimize the risk of prompt injections or, you know, data somehow, you know, exfiltrating out, or being inserted in mix.

George Lee: This leads me to something you've written about Marco, that I'm fascinated in how the managerial challenge will evolve over time, where we're not only managing humans, but we're managing AI agents. How will that shift the job of being a manager and what skills do we need to build to live in that world?

Marco Argenti: Yeah, that's an incredibly, you know, fascinating question. So, first of all, one of the challenges is to somehow inject call them, the cultural traits, the leadership principles, the tenets, like I like to call them, of

the organization into the AI agents the same way as you do with humans because cultures are very peculiar. Goldman is very peculiar. Amazon, where I was before, was very peculiar. And I tell you that people tend to kind of unify their decisions, their way of thinking around certain tenets, sometimes written, sometimes not written. And cultural traits, which, you will get when you're here. And they're hard to describe, but you see them when you see them. Whereas agents that you implement from the outside are going to be like employees that we hire the first day. And so the challenge is agents will get smarter and smarter, but not culturally smarter if you don't do something on that front.

And I think this is not something that has been talked about too much. We talked about mostly like specialization of expertise. We're not talking too much about in a way enhancement of the cultural aspect of agents, because those agents eventually will you know, even today we have, let's say AIs or kind of agentic AI that does things like providing, or creating drafts of research. For example, you know, in private wealth management, when it comes to portfolio analysis, to use agents, to summarize the nature of business of dozens of positions and in a way that puts a

Goldman lens into the answer. Okay? Even like something as simple as translation, which obviously should come out of the box. We have adapted the tool of translate AI to translate in a way that is actually, you know, very specific to our financial content, but also that somehow is specific of the Goldman language in a way, right? The way we call things, and you imagine just acronyms, how complex that is. And so those are just early examples. But in general, I think as we use agents to more and more, interact with people, but also produce content for people. I think there is a little bit of a sort of an extended example of the three laws of robotics from Isaac Asimov of there are the, you know, what are the actual leadership principles or tenets that an effective corporate agent should adhere to in order to fit into the culture of an organization. And I think that's a very complex, problem that, I don't think, has been solved yet, to be honest.

Allison Nathan: I mean, this hybrid workplace is just such a sea change from where we are today. I mean, do you see it really just driving corporate strategy overall?

Marco Argenti: Yeah. There are companies, like I was on a panel actually yesterday with a chief people officer of a

major software company. And her new title is chief people and AI officer because, they're already thinking that even today they want to look at how would we make sure that employees that communicate with our AI tools are not exposed to biases, for example. And also the cultural aspect of the answers is very important. And then agents are definitely something that is already knocking at the door. And so when you do like resources planning or when you actually like, you know, the training, right now is such a big deal because we are all pretty much AI illiterate. This is comparable to a change on pre- and post-computers. And by the way, funny enough, if you talk to people like in the, I don't know, forties and fifties or sixties about whether they could imagine a workplace without computers. A lot of people will tell you absolutely. You know, the computer is always going to be this big thing over there that a few people with white lab coats are going to use. But on a day by day, there's nothing like a physical signature or a paper. You're gonna sit down and read it and whatever, and think about like today, imagine, any company that doesn't have a computer, what would they do? They couldn't scale even like minimally. And I think in this new world, there is going to be actually an interesting- like you mentioned that, but

a bit of a generation, shift where, you now have this new generation, and my daughter is 20 and she can, knows how to use AI very naturally because she's been using it for, you know, at least two or three years. And a lot of the kids that are even coming into Goldman, they're extremely versed on AI. They're great prompt engineers and it took them years to perfect that. The same way as computers. And you know, like, I've seen that with, you know, with my parents or etcetera, that there was a point in time in history where if you actually never use the computer until you were like 40, you would be like never having ridden a bicycle until you were like 40 or 50, or never done certain things early enough that you will never be natural around it.

You could use it. But even today, when I look around, there's, I swear, there's people that are still typing with one finger. Or two. Okay. And those are fine. They can be very effective. But, there has definitely been a generational gap before on the use of computers and then on the use of the internet, and now with AI, I think is even more profound. And so the change happening in a generation that is not AI first is the same as a generation that was not digital first. And that's why I think that's one of the biggest challenges

for human resources, and human capital professionals to try to understand how to navigate into this. And what makes it, I think, different is that all this is happening in weeks versus years. And that is what I think reason some of the most profound questions on how to actually drive that change or how to actually somehow survive that change.

Allison Nathan: So how long will it take until we get to a hybrid workforce?

Marco Argenti: I think it's for developers, it's already kind of happening because there are today virtual developers that you can buy for you know, $20, $30 a month literally. And you can chat with and they do things for you. You're always going to need a developer because I don't believe in a future that is dominated by AIs and humans have a marginal, I mean, I, I tell you that that is not even in sight, not even for the most futuristic predictors. In fact, AI, I think, makes the role of humans even more important because you now have an added responsibility. Everybody becomes a manager. Everybody becomes a manager of agents. They can have hundreds if they want, and everybody needs to be somehow taking responsibility for

the work that is done. And the amplification of your good and your bad can be incredible. Your mistakes can be amplified and your successes can be amplified. And so you even have more responsibility. Therefore, this is a wake-up call to step up for everyone. And I think the centrality of professionalism and high judgment has never been more important than now.

George Lee: Speaking of being a AI native, you and I both push ourselves to be ardent adopters of the technology, but you have a very specific personal use case that I admire and I've seen a lot of cool evidence of. Why don't you talk about some of the ways that you use AI and the one in particular, you know, I'm thinking about.

Marco Argenti: So I do like in my spare time, which is not much, but I'm a little bit of a, you know, amateur musician. And so, I composed a few songs with, with AI that actually some of them, my band has actually performed live. And I have to say that AIs can do terrible compositions, but actually, if used properly, they really accelerate or make it, you know, make it much easier to kind of go through the creative process if you try to kind of help it for some of the mechanical part of composing a song

or like, you know, creating rhythms or loops or so there is a lot of musicians that are actually using it creatively. And, I'm not obviously one of them, but I aspire to be at least okay on that. And so I've been, I've been doing that with pretty good results. And all it takes is just knowing how to ask.

George Lee: First of all, Marco is being humble. He's a terrific musician, and I must say the music that he co-composes is remarkable. The lyrics are poignant. The rhythms are interesting. And so it's one of those things that you see that really makes you sit up and take notice of the progress of the technology and the ways that human creativity can be used, can be elevated with it. So Marco, thank you so much for being with us. Fascinating discussion. We could go for hours.

Alison Nathan: Yes, absolutely. I mean, we started this conversation saying adoption, we got to give ourselves a break. We're only a year and a half into this. So I'm really interested to see how this all evolves. Thank you so much, Marco.

Marco Argenti: Thank you so much. Thank you.

Allison Nathan: And thanks to you George.

George Lee: Great, that was fun!

Allison Nathan: This episode of Goldman Sachs Exchanges was recorded on Thursday, March 6th. I'm Allison Nathan. Thanks for listening.

The opinions and views expressed in this program may not necessarily reflect the institutional views of Goldman Sachs or its affiliates. This program 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. Each name of a third-party organization mentioned in this program 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. The content of this program does not constitute a recommendation from any Goldman Sachs entity to the recipient, and is provided for informational purposes only. Goldman Sachs is not providing any financial, economic, legal, investment, accounting, or tax advice through this program or to its recipient. Certain information contained

in this program constitutes “forward-looking statements”, and there is no guarantee that these results will be achieved. Goldman Sachs has no obligation to provide updates or changes to the information in this program. Past performance does not guarantee future results, which may vary. Neither Goldman Sachs nor any of its affiliates makes any representation or warranty, express or implied, as to the accuracy or completeness of the statements or any information contained in this program and any liability therefore; including in respect of direct, indirect, or consequential loss or damage is expressly disclaimed.

This transcript should not be copied, distributed, published, or reproduced, in whole or in part, or disclosed by any recipient to any other person. The information contained in this transcript does not constitute a recommendation from any Goldman Sachs entity to the recipient. Neither Goldman Sachs nor any of its affiliates makes any representation or warranty, express or implied, as to the accuracy or completeness of the statements or any information contained in this transcript and any liability therefor (including in respect of direct, indirect, or

consequential loss or damage) are expressly disclaimed. The views expressed in this transcript are not necessarily those of Goldman Sachs, and Goldman Sachs is not providing any financial, economic, legal, accounting, or tax advice or recommendations in this transcript. In addition, the receipt of this transcript by any recipient is not to be taken as constituting the giving of investment advice by Goldman Sachs to that recipient, nor to constitute such person a client of any Goldman Sachs entity. This transcript is provided in conjunction with the associated video/audio content for convenience. The content of this transcript may differ from the associated video/audio, please consult the original content as the definitive source. Goldman Sachs is not responsible for any errors in the transcript.