



The Economic Club of New York

119th Year
839th Meeting

Sebastian Mallaby
Paul A. Volcker Senior Fellow for International Economics
Council on Foreign Relations

April 1, 2026

In-Person/Hybrid Event

Moderator: Gillian Tett
Editorial Board, *Financial Times*
Provost, King's College, Cambridge

Introduction

Chair Robert K. Steel

Welcome everyone to the 839th, yes, 839th meeting of The Economic Club of New York. My name is Bob Steel, and it's great to welcome you here on behalf of the Club. Let me also extend a warm welcome to the students that are joining us today virtually from Harvard Business School and Duquesne University.

The Economic Club of New York is proud to stand as the nation's leading nonpartisan platform for discussion on economic, social, and political matters. For more than a century, the Club has hosted over 1,000 preeminent guest speakers contributing to this tradition of excellence over a long period of time.

We continue that today with part of our Author Series, where we're really honored in an incredibly timely way to have Sebastian Mallaby here to talk about his new book.

Sebastian is the author of several previous books, including bestselling *More Money than God*, a former *Financial Times* contributing editor, and two-time, two-time Pulitzer Prize finalist. Mallaby is the Paul A. Volcker Senior Fellow for International Economics at the Council on Foreign Relations. Paul was a member of our Club and spoke regularly so it's great to have that connection with Sebastian here also.

Today's discussion is part of our planned series of important authors and today we're going to explore Sebastian's latest book, *The Infinity Machine: Demis Hassabis, DeepMind, and the Quest for Superintelligence*. We were talking before and, you know, we started this a handful of years ago and to be so prescient as to imagine what a great topic would be and giving us the privilege to engage with him on this topic at this time when he began several years ago, we're so very fortunate.

Today's format will be a conversation moderated by Gillian Tett, who we're really honored to have here. Gillian is a great friend. She's a columnist on the Editorial Board of the *Financial Times* and also provost of King's College at Cambridge. What we'll do is have a conversation with them and then if we have time, we'll ask some questions at the end depending on how time runs, but we will be ending promptly at one. So now let me conclude by inviting them to the stage and we'll begin the program. Thank you.

Conversation with Sebastian Mallaby

GILLIAN TETT: Well, hello everybody. It's a terrific honor to be here at the University Club which I do think is one of the most beautiful clubs in New York and also dealing with The Economic Club of New York that I've dealt with a lot over the years and have great admiration for.

But what's even better is being here with an old friend, Sebastian Mallaby, who, as you've heard, has been a journalist and commentator for many, many years, and also an author. And just in case you think that this is an official British takeover, I can assure you that although Sebastian and I both sound British, in fact, we're both dual nationalities and we're both British and American. So you can see us through whatever lens you want today or both ways.

But we're talking today about someone who has told me many times that not only is he proud to be British, but he has resisted attempts to move him to America many times. He insists on staying in London even though Silicon Valley has been obviously the main focus of attention in AI in recent years. And we'll come onto the question of whether there are going to be any challenges in Silicon Valley in the future. But that person, of course, is Demis Hassabis, who was indeed at Cambridge University where I am, as you just heard, head of King's College. Demis himself was actually at Queen's College just next door. And he was there as an undergraduate and then he went on to do his PhD in London and then went on to great things.

He is, in some ways, one of the most fascinating figures today in AI. Because although he hasn't grabbed the spotlight like Sam Altman, in many ways his influence has been – one might argue – more significant. And most importantly of all, he is forging a path that is self-consciously quite different from Sam Altman and is very strongly committed in his

eyes to trying to unleash AI for good, even as he tries to take AI to its extremes in terms of artificial general intelligence. In some ways, he reminds us a lot of Robert Oppenheimer in terms of his vision of how nuclear energy could be used for good, but then of course it was later turned for bad as well.

So perhaps I can start by asking you, Sebastian, this is a terrific book. When did you decide to start writing it? Because Demis wasn't exactly a household name when you started all those years ago, was he?

SEBASTIAN MALLABY: Yes, well, I had the idea in the summer of 2022. And at the time, if you were following DeepMind and, you know, I think we both probably knew him from tech conferences and so forth, you would show up at a tech conference in Europe and there would be this diminutive, smiling, grinning, sort of boyish figure with, at the time, hair falling down into his eyes a bit. And, you know, he would get on the stage smiling disarmingly and then suddenly start riffing about Immanuel Kant, the nature of intelligence, the connections between neuroscience and computer science and all these incredible ideas, and I always found him captivating. So this is one impulse to try and do a book about him.

GILLIAN TETT: He's one of the most unlikely billionaires out there. I mean a nice billionaire.

SEBASTIAN MALLABY: So I thought he would always be, he would be a great lens, a lens probably at the time into a niche topic, but one that I found interesting and that seemed to be gathering a bit of momentum. I thought it would break out eventually. And so I went to see him in November of 2022, and I pitched him on the idea of giving up so much of his time that I could really write a good book, hours and hours and hours. And, you know, I said to him, look if it's true as you say in your speeches that AI is going to be the most important invention in human history, it follows, Demis, that you will be one of the most important people in human history. And therefore, you don't have a choice, there's going to be a book about you.

And anyway, if you're going to inflict on the world a technology that changes how people bring up their children, how they operate at work, how they think about themselves as humans, because there's now a rival source of machine cognition challenging what we thought was our identity, you can't do that to the world without explaining your motivations, your reasons for doing it.

And so not only will have a book about you, you should want to have a book about you because you can't get away with just not explaining yourself. So that was the pitch in November 2022. He seemed interested. I thought it would go mainstream eventually. And then November 30th of that year, ChatGPT comes out and it goes from the fringe to the mainstream way faster than I had ever imagined.

GILLIAN TETT: So before we start talking about some of the shifts in AI technology and what that means because that's very important, just give us a quick sense of where Demis came from and what kind of person he is because it's a kind of background which is not familiar to most people in this room. Ironically, it's actually a part of London where I myself grew up and am very familiar with it. Just give us a sense of what it meant to be sort of an immigrant, or rather the child of immigrants in North London.

SEBASTIAN MALLABY: Yes, so Demis's father was of Greek Cypriot extraction and his mother was Singaporean Chinese. She had been an orphan for part of her childhood in Singapore. So we're talking about, you know, moderate income immigrant couple in North London sending their son to the local government school. And the whole thing was a sort of, in a way, a typical London story, right?

GILLIAN TETT: Yes. North London.

SEBASTIAN MALLABY: I think about four-fifths of the workforce in London is not born in London. And so maybe to an extent that people in the U.S. don't recognize, you know, London is quite a melting pot.

GILLIAN TETT: It's a complete melting pot. My kids' school, when they were young, there were literally two parents out of the entire class who were both White British, and

that's normal. No one pays any attention to it at all.

SEBASTIAN MALLABY: And, by the way, Demis's co-founder, much later on in the story, Mustafa Suleyman, was also from North London. His father was an immigrant Syrian cab driver who couldn't really speak English very well. And there were no books at home, no music was allowed. It was a very devout Muslim family. His mother was a convert to Islam. So again another immigrant story from North London who then became the co-founder of DeepMind. So, yes, it's a very special kind of DNA.

And Demis, you know, escaped or kind of rose up out of that background just through a shared force of intellect. He got up on a chair when he was four years old, watched his father playing chess and looked at the board and very quickly just understood the game. And when he was 10, he had become the top youth chess player in Britain and the second best in the world. He just was that clever.

And he'd had a very ordered didactic childhood because, you know, his parents weren't really equipped to lead him through the reading of science and novels and his magpie interests. That was very much him by himself with books. He dropped out of school for a whole year because he was so focused on chess and just taught himself in his bedroom what he needed to know. So kind of isolated, super competitive because the chess circuit was brutal. You know, the kids who played, they used to put these wooden slats

under the table to prevent one player from kicking the other one in the shins to distract them. So very, very competitive, very self-taught, and very self-made.

GILLIAN TETT: Yes, it's ironic because America is famously a land of immigrants, Silicon Valley being built by immigrants. People often ignore the fact that in a place like the U.K., there is this very dynamic, you know, hardscrabble immigrant culture in places like London as well, which in some ways is more of a melting pot than New York ironically. So he comes out of that, he gets himself to Cambridge, to Queen's College. Sadly, not King's. We would have loved to have claimed him as one of our own, but we didn't.

SEBASTIAN MALLABY: You did have Keynes, so it's not too bad.

GILLIAN TETT: We did have Keynes, and we also had Alan Turing and Geoff Hinton, so we can't have everybody. So he basically goes, he gets through his undergraduate degree and then goes to UCL for computer science. How does he then become an entrepreneur?

SEBASTIAN MALLABY: Well, again this is totally extraordinary. I mean in 1997 when he left Cambridge, leaving Cambridge and starting your own company was just not a thing.

GILLIAN TETT: Absolutely not. And, by the way, we're now trying to change that, but back then entrepreneur was a dirty word, and not just because it's French.

SEBASTIAN MALLABY: Right, right. And there's a story I tell about how he thought he would found this company initially to build video games, but with the ultimate goal of building AI. He already had this ambition to be the man, the person who brought AI to the world, and that was kind of almost a messianic vision from a very early period in his life. And so he decided to get that by starting a company, and he went off to raise money from people who sort of called themselves venture capitalists, but they weren't really VCs in the way that anybody in this room would recognize. And the upshot of the first meeting which took place over far too much alcohol at lunchtime...

GILLIAN TETT: Which, by the way, is another old British tradition.

SEBASTIAN MALLABY: ...the upshot was that really a banker, not a VC, said to him, you know, oh, I love your chess stories, but never mind your silly company. Why don't you come and work for my bank, you'd be a great currency trader. So that was how much, even a VC in London, was supportive of actually starting a company. So, yes, totally out of the ordinary.

GILLIAN TETT: So then he gets the company against the odds. It's pretty much self-

funded initially, isn't it? I mean it's bootstrapped up.

SEBASTIAN MALLABY: He bootstraps it. He had made some money. You know, as a 17-year-old he had written the code and come up with a design for a very successful video game and it had sold like more than five million copies. And the boss of the gaming company was so amazed by Demis's creativity that he tried to bribe him not to go to university. He offered him more than a million dollars in today's money not to go. And Demis didn't cash the check because he wanted to go and study. Again, in Stanford probably the kid would have dropped out.

GILLIAN TETT: And he would be funded by the likes of Peter Thiel straight away.

SEBASTIAN MALLABY: Right. In London, amazingly the kid, even though he had parents who were of modest means, said no to the money and yes to the science and studied computer science. So he had some money behind and he bootstrapped this post-Cambridge company, Elixir, and it did okay. He did produce some games. They were kind of moderately successful.

But interestingly, and this is a classic for people interested in how startups work, he blew it up by being too ambitious. He wasn't willing to kind of balance the vision with the pragmatism. It's great to have a vision. It's great to tell the premature truth sometimes.

It's great to be super ambitious. But if you actually have an ambition that can't be physically built in the Silicon that you have available at the time, you know, you're going to fall flat on your face and that's sort of what happened to that first company.

GILLIAN TETT: So like many entrepreneurs, you know, he had a failure, a qualified failure. So what then started him onto the path of success?

SEBASTIAN MALLABY: Well, he went off and kind of recharged his batteries because he was totally burnt out. And his idea of recharging his batteries was (a) to get married to his Cambridge girlfriend. Maybe that's sort of normal. But then secondly, on his honeymoon to take a very large pile of neuroscience papers about memory and imagination and enrolled in a PhD in neuroscience. So that was his next move. And so to recharge he writes papers about memory and imagination in the hippocampus and wrote one of the most cited neuroscience papers of the decade in the process.

And the point of doing neuroscience is to understand human intelligence as an existence proof for artificial intelligence. If you're going to build AI, you needed some intuition of what even is intelligence, right? And then it was coming out of that PhD and doing a post doc at University College London with some stints at MIT and Harvard. Of course, being Demis, he didn't just do MIT or Harvard. He did both simultaneously. One scientist friend of mine who read the manuscript said that's impossible, you're making

that up. So I messaged his supervisor at MIT, Tomaso Poggio, who was Demis's supervisor. Is it really true that he also had a post-doc at Harvard at the same time?

Yes, it was true. It was Demis. It was true.

So he comes out of this post-doc and in the process of the post-doc he meets Shane Legg. And Shane Legg is a New Zealand, another immigrant to London, a New Zealander, who is a computer scientist who has sort of done this PhD on the definition of artificial general intelligence and is the only other person in the whole of this post-doc program at the Gatsby Unit at University College London – the only other person who was thinking in terms of ambitious AI. Not just let's build a kind of toy experiment that can maybe, you know, recognize a rough picture, but let's really imagine machines that are cleverer than humans. In all of this academic institution, there were two people who believed that that could be done and that was Demis and Shane. So they linked up and that was the foundation of DeepMind.

GILLIAN TETT: And then Mustafa enters the picture, doesn't he? Mustafa Suleyman, the other immigrant in London who also had this extraordinary vision, and in some ways is quite similar to Demis.

SEBASTIAN MALLABY: Yes, I mean this is something where, you know, Demis has tried to, and he ended up firing Mustafa from DeepMind after nine years of working with

him. And he really didn't want me to kind of write about this or excavate what happened. And so he would say, you know, first of all, you would go to DeepMind events and there would be a story told about DeepMind founded by the two co-founders, Shane Legg and Demis, and Mustafa wouldn't even be mentioned. And, you know, then he would say, well, yeah, Mustafa, he was there for a while. He was my brother's friend.

And so, of course, I went to see Mustafa and I spent almost 20 hours with him. And the relationship was actually quite deep because Mustafa was the best friend of Demis's younger brother and would go and stay – Mustafa, you know, I told you about the immigrant Syrian cab driver, what I didn't say earlier is that the immigrant cab driver left London to go back to live in Syria to find a new wife when he got separated from Mustafa's mother. And so both the parents left leaving the 15-year-old Mustafa to look after himself and his 14-year-old brother and somehow not starve.

And so Mustafa would make friends with people. He was extremely bright and kind of good at connecting. And he made friends with this guy called George Hassabis who was Demis's younger brother. And because Mustafa had nowhere to live, he would sleep in the spare room at George's home which was the vacated room of Demis. And so there he was sleeping in Demis's room and Demis would come over to see the family sometimes and Mustafa would strike up a conversation when he could.

And they went into business together doing real estate. Demis had made a bit of money from his first ___ and Mustafa organized the investment of it in real estate. And they connected and when Demis needed somebody to kind of write the business plan, find the first office, sort out the administrative side of setting up a business, he brought Mustafa in as the third co-founder.

GILLIAN TETT: Right. Like many startup stories, there's a lot of complicated personal history in all of this and what happened with Demis and Mustafa was, in some ways, is tragic. I mean they remained rivals, frenemies. They're both playing this important role in shaping AI.

So let's turn to the question of AI because they finally create DeepMind. It's not particularly in the headlights at all. It's sort of, you know, just trundling on doing this research, but bit by bit they begin to create some presence. They take slightly different paths though, don't they, to ChatGPT and Open AI in terms of how they imagine AI. And it's important to stress that because some people think AI is just one thing. There are actually different variants of AI. So tell us what they actually do about AI and how it's different from what Sam Altman did.

SEBASTIAN MALLABY: Sure. Well, the first thing to say is, you know, they founded DeepMind in 2010 and this is before the ImageNet breakthrough where AI could finally

recognize a cat. So in 2010, nothing worked in AI, so it was pretty gutsy to even have the idea of a company. So, you know, they managed to talk money out of Peter Thiel. They kept going. Nothing works at first, but they do have an idea.

And the idea is really that two strands of AI, separate computer science traditions, are both maturing, and if you put them together the two could do amazing things. And the two strands, I think this is sort of helpful for people to understand just as a very, you know, there's AI, but if we go one level deeper what's in the AI? I think the kind of clarifying distinction is that there's on the one hand, deep learning. And deep learning is a system that takes in tons of data so, you know, it inhales the internet, it trains on that data, and it can start to recognize patterns in the data. That's deep learning.

Then the other thing is reinforcement learning. Reinforcement learning is learning through experience, through trial and error, through interaction with an environment. So you have a system that, for example, is set up to play the game of Atari, and the most simple example would be Pong, it's like a ping-pong game. And the computer has one paddle and it hits the ball back to you and you're the player and you're controlling this paddle, moving it up and down to try and block the ball. In fact, in my previous book about venture capital, I wrote about Atari because it was the first investment by Sequoia. And I joked that, you know, the system is so simple, the game is so simple so that even the most inebriated bar goer could play it because that was the market for the

game at the time.

So they begin with a simple game and you've got to train the paddle to, you know, the AI is going to move the paddle up and down to block the ball. And what it does is it moves at random at first, and then if it, by total coincidence hits the ball, it says, aha, whatever I just did, that was smart, I should do more of it. And so by getting signals of reinforcement from when something goes right, the system gradually builds out the picture of which action is going to lead it to succeed in its task.

And so in the Atari example, you have deep learning, which is kind of ingesting all, both the picture of the game and so making sense of the picture of the raw pixels and turning that into a mathematical representation of what the board is. And so you need some deep learning, but you also need reinforcement learning to just try out, trial and error. And as the story of AI goes forward, at different points reinforcement learning is the dominant source of progress. At other times, deep learning is the dominant source of progress. And DeepMind was really betting that you needed this.

GILLIAN TETT: Right. What DeepMind did not do, though, was pay much attention to the issue of language because we now live in an era where Silicon Valley is betting the farm on large language models, LLMs. In fact, that's not the only type of AI out there and they may yet be overtaken by other types of AI, and I suspect that's going to

happen. But pre-the ChatGPT moment, essentially people like Demis were kind of ignoring language, weren't they?

SEBASTIAN MALLABY: Yes, or at least they were sort of not, I mean they refused to stress it. And this was a very central part of my discussions with Demis when I was doing the research. You know, we would go meet in a pub in North London near his home and we would climb a sort of musty staircase at the back of the pub, which was sort of a bit obscured. Nobody else knew about it so it was great because then it was just two of us up in the upstairs room under a kind of, you know, dusty old chandelier.

GILLIAN TETT: If you had been in Silicon Valley, you would have either been in a wine bar or you'd be sitting around a campfire with ___ or something like that. But in London we do pubs.

SEBASTIAN MALLABY: We do pubs. And, you know, a very fine pub it was too because of all the privacy it gave us. And so I'd be sitting there for two hours asking him questions, and he would riff. And one of the questions that came up all the time because this was against the background of ChatGPT and Open AI becoming the leading most famous AI lab because of ChatGPT. I'd say, look, Demis, you know, you started before ChatGPT. You built all these great systems, the AlphaGo system that became the world champion at Go. You had AlphaFold, which, you know, unraveled all

the protein shapes in nature. And if AI was you, AI was DeepMind, why did you let these guys overtake you? What happened?

And the answer came down to language as you rightly said, Gillian. So the answer came down to the idea that, look, language is a system of symbols that represent things in the real world. Pictures are also symbols of things in the real world. We could develop AI a long way so that it matched the pictures and the symbols together, you know, the picture of the cat and the word cat.

But in the end, to be really intelligent you need to actually interact with the physical environment. You need to know if I pick this up, I feel weight. What is weight? If you just downloaded all of Wikipedia and you read all of Wikipedia, Demis said to me, would you understand weight? Would you understand that if you tip it, the water will come out because of gravity, that if you drop it might shatter? These are all things, and I'm really kind of repeating exactly what he did to me at the pub. He picked a glass of water up.

And therefore, just trying to build a super intelligent machine by maxing out on language modeling is going to be a dead end. This was his view. And he maintained this view through 2018 when the first GPT model came out.

GILLIAN TETT: The crucial papers came out basically creating the basis for the

Transformers papers.

SEBASTIAN MALLABY: Sure. And then 2019, the same thing. So 2017 was the Transformer model paper. That was from Google. 2018, Open AI has jumped on the Transformer paper and created the first language model, the first GPT. 2019, there's a better version that comes out. 2020, there's Version 3. And then suddenly Demis says, oh, my God, I was wrong.

And what he realized was there is so much information in language that if you download all of Wikipedia, actually you do know what weight is because it's described in so many ways. And if you download literature and so forth, you end up being more grounded, or at least be able to simulate groundedness, than he had predicted. And so then began, after 2020, one of the most epic sort of catch-up business race stories of all time.

GILLIAN TETT: Right. I mean what happened, that was astonishing. And, by the way, what happened might happen again in the future with AI. I mean in many ways this is an echo of what, or a preview of what could happen next time around with AI as the current models get superseded. But would you quickly say, in between all this, before Demis launches his war with Open AI, he had already, and then went on to develop this extraordinary AlphaFold platform, which I don't know whether everyone in the room is familiar with that, but that has revolutionized parts of life science. I mean I see it every

day in Cambridge how it's completely transformed the way that life science research is being done. And that's what won him the Nobel Prize. Tell us briefly what that is and why it matters, because that's kind of a different branch of AI, isn't it?

SEBASTIAN MALLABY: Yes, when Demis was at Cambridge, he had this biologist friend who made a sort of passing comment to him about how there's a longstanding grand challenge in biology. A Nobel Laureate from the 1970s called Anfinsen gave a lecture saying, look, you know, much of nature is composed of proteins. These are tiny, intricately shaped, mini bio-factories that do a lot of the things that we need to have done in our bodies, also in plants and nature and so forth, and they are made up of amino acids strands. And in the amino acid strands there is a genetic code and depending on the sequencing of what's in the strand of that genetic code, the code, said Anfinsen in the 1970s, the code predicts the structure.

So if you stretched out the amino acid strand and you read the code, the genetic code, you could predict that when you let go, it would act like a self-executing origami model, twist itself up into this extraordinary intricate shape. And that shape could be predicted if you decipher the code in the strand. And this was just posited in the 1970s. And in the 1990s when Demis was at Cambridge, his friend said that would be cool to solve that. So Demis just stored this in the back of his brain.

And then there was a moment in 2016 when DeepMind defeated the Go champion in South Korea. And as Demis walked out of the Go, sort of, you know, the kind of exhibition hall where the Go match took place, he was walking down the street in Seoul with his friend from Cambridge, David Silver, who also was the leading reinforcement learning scientist at DeepMind. And they were chatting about having just won Go and rather than rest on his laurels, Demis said, you know that thing I talked about in Cambridge, protein folding, I think we can do that next.

And he then set a team to work and just as in Go, if you think about Go, Go has 19 squares on one side and 19 on the other. So the first move, there are 361 options when you put your counter down. Now that one square is taken so the next move there's 360 options. So you do a factorial and you quickly get to some infinitely large number of combinatorial possibilities about how Go might work. By the way, that's why I call my book, *The Infinity Machine*, because AI is a system for unraveling meaning from an infinity of possibilities. So if you could solve Go with a massive combinatorial complexity, why couldn't you solve protein folding, which is another problem where that strand of DNA can fold itself up into an origami with, you know, billions and billions upon billions of different folding options.

So in computer science terms, one success was the springboard for the second success. And sure enough, four years after that moment in 2016 when he formulated

the ambition, his team at DeepMind had managed to predict all protein shapes in nature just from the DNA code. And that obsoleted the other way of finding out the shapes of proteins, which is called X-ray proteinography, and I won't go into all of that. But it would basically take a PhD biologist most of the PhD, like five years, to piece together the shape of one protein. Once DeepMind cracked the code, it could do all the proteins in nature in a matter of a couple of weeks.

GILLIAN TETT: Which if you're looking for an example of how AI can do good and bring about real benefit, frankly I don't know many other examples which are quite so powerful, both in terms of the scale of the transformation in life sciences and the potential for good. I mean it tends to get overlooked what they did, but it's astonishing. However, it did get overlooked because then what happened of course was, we moved into the era of transformers, of ChatGPT, DeepMind gets bought by Google. Open AI comes out, grabs all of our attentions by releasing ChatGPT and overnight AI becomes a sort of household name and a sensation. Demis realizes he's fallen behind. What does he do next? And can he catch up?

SEBASTIAN MALLABY: Yes, well, I mean I went to see him right after ChatGPT came out. I had pitched him on the idea before it came out. Then it came out. Then I had to go see four of his advisers to check me out and read my past books. And then I was summoned for the next meeting which was a dinner. And I think it was at that meeting

or the one after, he said to me those guys at Open AI, they have parked their tanks on our lawn, on our front yard. This is war. And, you know, you could see that super-competitive chess playing glint in the eye. And you knew that he was going to try to do a comeback.

Now, at the time, when I went to Silicon Valley to see my friends in the venture capital world and I talked about the large language model race, they would say, look, game over. You know, Open AI went first. It's now gone viral faster than any other product in the entirety of technology history. No way anybody else can catch up. Forget it, you're writing a book about the wrong person. And I would say I'm not so sure. Why do you think it's impossible to make up the gap? And they would say, mindshare, whatever that means? You know, basically it meant that Sam Altman was the best at creating buzz for his products and that he's the ultimate Silicon Valley networker. And they were part of the network and so they perceived some of them to be unbeatable. And I wasn't convinced.

And Demis set about, you know, a few things happened. First of all, they merged the AI team that Google had in Mountain View, which was called Google Brain, with the DeepMind team in London. So now you had a 5,000-strong AI bench. And they put Demis in charge of both. That was kind of a strange move in the sense that, you know, Demis was based in London, not in Google headquarters.

GILLIAN TETT: And refused to move to Silicon Valley.

SEBASTIAN MALLABY: Completely refused to move to Silicon Valley. You know, Jeff Dean, who had been the head of Google Brain was a much more longstanding Google headquarters insider. But I think Sundar Pichai had always understood Demis's exceptional ability and always knew that when the moment came when AI was going to become the central engine of Google's future, Demis was going to be the person to drive that, and so he put Demis in charge.

And any business school professor would have said, listen, merging two different laboratories is really difficult. Merging them in the middle of a, you know, existential race against your rival, Open AI, is even harder. And merging them when they are eight time zones apart from each other, California to London, is almost impossible. And particularly if in the past, as is the case with these two labs, they've been daggers drawn because it used to be that Google Brain would produce an AI thing and DeepMind would produce one and then they would both want to train it on as much computational resources as possible. And so they hated each other because they were always fighting about who got the computational power. So it was not likely to work. And yet two and a half years later, it had worked. By the end of 2025, Demis's Gemini model was ranking higher than the Open AI one.

So I think it is a kind of historic business case study of how you do a merger and make it work. And I would say a final thing on this, the relationship between Sundar Pichai who was often underestimated and Demis Hassabis is the single most important buddy act in global business today. Right? Those two people, because Sundar runs the company, delivers the resources to Demis so he can do the scientific thing. They both need each other and it really works.

GILLIAN TETT: So that raises two questions, and I'm then going to turn to the audience for questions, so start thinking some great questions now. But the first question is, you know, I've asked Demis many times why he won't move to Silicon Valley. And he's absolutely obstinate about this and nobody is going to last a launch in his obstinacy. But he keeps insisting that by staying out of Silicon Valley in some ways he has a benefit of a different perspective. Silicon Valley is too inward looking. It's too much of a uniform perspective, etc. etc. Do you think that's correct? Do you think part of the secret sauce is not being in Silicon Valley? And do you think that, you know, Europe/London can ever compete on the stage in this area?

SEBASTIAN MALLABY: I think I have a split answer on this one, on whether London can compete, I do think it can. I do think it's underestimated. I think actually one of your colleagues, John Thornhill, just had a column in the *Financial Times*, which I agreed with, which said that, you know, actually quite a lot of American venture capital dollars

go into European startups.

And I go to startup conferences in both Silicon Valley and in Europe and I'd say that, you know, the difference has shrunk. You go to Europe, and there are lots of people on their second company, their third company, they're super-ambitious about that, how they're going to build it. Most of them have spent some time working in Silicon Valley and they understand that DNA. They have financial backers who are either American VCs or American-trained European VCs. So I really think, you know, on the micro, Europe is way more promising than people give it credit for. On the macro, if you look at government debt, slow growth, all that stuff, too much regulation, I'm quite pessimistic about Europe. But on the micro, I'm positive.

I wouldn't go so far as to say, though, that DeepMind benefits from being in London. I mean I think, you know, there's a story that there's a lot of talent in Europe scientifically, and that's true, and that you can hire it into London. Yes, that's true. But I think quite a lot of talent is ____ to America. And so I think it's pushing it a bit too much to say it's better to be in London, but I think it's fine to be in London and DeepMind is proof of that. And, by the way, so is the, you know, the first – I think I'm right – the first Covid vaccine to be approved by the FDA was the Pfizer-BioNTech one invented by...

GILLIAN TETT: In Oxford...

SEBASTIAN MALLABY: No, that was the AstraZeneca. AstraZeneca...

GILLIAN TETT: Invented the Pfizer one with a metabiologist, by the Turkish...

SEBASTIAN MALLABY: Yes, exactly, a Turkish immigrant to Germany, another immigrant story.

GILLIAN TETT: Absolutely. Immigrants, they get the job done. But in terms, by the way, I would echo that the shift in attitude towards entrepreneurship in London in the last couple of decades has been remarkable. I mean, you know, King's College has created the first Entrepreneurship Lab over Oxford or Cambridge College. It's going absolute gang-busters. And these days half the kids want to be entrepreneurs, which is a total sea change from 15, 20 years ago.

Last question before I turn it to the audience is that you have a quote from Demis in the book about his desire to basically go into the next stage of AI, which is artificial general intelligence. That would be just LLMs, it's much wider than LLMs in a much bigger sensory way. And you asked him why he wants to do this and he says, "the true reason to build artificial intelligence is to draw closer to what might be called God – the intelligence that may presumably have designed everything around us." And what comes out in the book is that he really believes, firstly, that AGI is possible, that he

wants to get there and that it will be good. And that his role in history is to unleash good AGI, safe AGI that will benefit mankind.

I've heard both Elon Musk and Sam Altman express sort of variants of that theme, but never in a way that feels so deeply heartfelt in terms of its good intentions. Maybe I'm being unfair to Altman and Musk but, you know, you can argue about that. So the question is, firstly, is it realistic that he will take us to AGI? And secondly, is he going to end up looking like Oppenheimer in the idea that he's unleashing a technology that he thinks is going to be good but ends up having actually terrible consequences?

SEBASTIAN MALLABY: Well, I think, you know, reaching AGI is almost in the bag. I mean I think it's definitely going to happen. There's fuzziness on the precise definition of what we mean by that. But we already have, you know, if you look at Gemini or pick your favorite large language model, you know, Claude or ChatGPT, these things are artificial, they are general, and they are intelligent. So, you know, we're not far off. And furthermore, the code is now generating code. So if you ask Demis or any of the other lab leaders how much of the next iteration of your model is coded by the existing model, it's a high percentage. So that recursive self-improvement loop is here.

GILLIAN TETT: Otherwise, there's reproduction.

SEBASTIAN MALLABY: Yes.

GILLIAN TETT: But not in a biological sense.

SEBASTIAN MALLABY: What we haven't got yet is quite, we haven't got, you know, the machines being sort of autonomously setting their own objectives and having a survival instinct. I mean humans have evolved over centuries to survive, to fight to survive, to pass on our DNA. Machines don't have DNA. They don't have that intrinsic desire to survive, but I think that's sort of emergent, which is worrisome. And that's why AI safety is such a huge deal.

If you imagine machines which are much more intelligent than us, if they don't fight to survive, that's not threatening. If they do fight to survive, they may fight us to survive and then we could be in trouble. So I think that is where we're headed. I think governments should be way more proactive in controlling the nature of this emergent survival instinct. Some people would call that consciousness, but I think survival instinct is a more specific description of what we should be worried about.

And so I think it's possible, therefore, yes, that Demis will bring this into being and that it could be dangerous and bad. And yet, he has wanted, ever since he met Shane Legg in 2009 and had the idea for DeepMind, he's wanted to make it beneficial and good. And

that was what AlphaFold was about as you correctly said. And so he is a case study in this sort of tingling sense of what does it feel like to have your hands on the 21st century version of the nuclear material, to have the responsibility for bringing this into the world? How do you sleep at night? How do you process that? And in some ways that was my deepest reason to want to write this book is that, you know, a portrait of somebody in such a consequential position is to me deeply fascinating.

GILLIAN TETT: And if you had to choose between Musk, Altman, and Hassabis, who do you most admire? And who do you most like?

SEBASTIAN MALLABY: Oh, I think it's easy to say Demis Hassabis by far. I mean, you know, I think he really has good values. He really does want to make AI good. When you write these books, and we discussed it beforehand, you know, you really stress-test somebody's personality. Because, you know, Demis began by speaking to me for 30 hours in that pub and then I lost count but there was a further period of, you know, I had promised to check quotations with him, so we did that. And in the process, you do that with some people and they will go ballistic and you see a different side of their personality.

And with Demis, of course, you know, his whole reputation was at stake in what I said. So it wasn't a totally easy interaction and he did get his lawyer to hound me about one

chapter about trying to stiff Google and spin out, which he secretly did and people don't know that story, but it's in the book.

GILLIAN TETT: It's a great chapter.

SEBASTIAN MALLABY: So it wasn't an easy discussion but he was always reasonable. He didn't go crazy, and I've dealt with many people who have gone crazy. So I think he is a good person, and I have no hesitation in saying that I would like him to be the one who had this sort of most influence in how AI turns out.

GILLIAN TETT: Well, anyone who is humble enough to meet in a pub to talk about those grandiose issues is probably a good bet for the future. But anyway, so we have a chance for some questions. So if anyone would like to ask a question, raise your hands. It will be courteous but not compulsory to briefly identify yourself and please keep them short. Let's do one, two, three...Frank, why don't you start...

QUESTION: Frank Brosens, Taconic Capital. You talk about guardrails for AI as being necessary. I think when Demis sold his company to Google in 2014, he insisted on an Ethics Board. One could argue that for these 12 years, that hasn't really slowed him down. But as AI gets more dangerous, is it likely to slow him down in the future and affect his ability to actually merge his...

GILLIAN TETT: So are Ethics Boards a bad idea?

SEBASTIAN MALLABY: You know they did try very hard to have an Ethics Board and that, in fact, was the fight with Google that I sort of reported out in my book. And they, you know, they had these big ambitions. They wanted Barack Obama to be sitting on it. They wanted sort of really credible independent people to have the final say over how AGI is deployed. And it didn't work.

In the end, Google decided, and one can understand why, that empowering independent people to have a say over core things that go on inside your company is probably a violation of your fiduciary duty. And it's very difficult to get people in those positions who don't then take the opportunity to essentially protect their own reputations by leaning in the direction of being critical of a company even when it kind of does the right thing.

And there's a whole story I tell which is on a smaller level about AI for healthcare which DeepMind tried earnestly to deliver. And there was a smaller oversight body to oversee its health work and instead of kind of properly judging it fairly, this group couldn't resist dinging them because they had an American parent company and, you know, Google must be evil because it's American. And so it really taught Google a lesson in the danger of empowering a board and so they got cold feet and they pulled out of these

negotiations.

And so I think Ethics Boards turn out to be super hard to work. You know, Open AI did it in a different way. They began as a nonprofit and tried to graft on a full profit. And we all know the story that when the nonprofit board fired Sam Altman and that lasted for about five days and then Altman fired them. So it's just very hard to make these things stick.

I think that, you know, in the end Demis's position is you have a collective action problem. Multiple labs in multiple countries racing to deliver AI. If one lab tries to be safe and careful and the others don't, it doesn't really make the world any safer. So that's not a good call. Neither for yourself and your employer, nor for the world. And Demis could also say, well, I'm not making this safe. I can't make it safe. I'm going to take a professorship at Princeton and go and think big thoughts. But that wouldn't make the world better either, right? So I think his dilemma is that he wants it to be safe. That really requires governmental intervention to enforce safety standards on everybody.

And what I think he's doing is lobbying behind the scenes for governments to see that and take action. The problem is the Trump administration is just completely not interested. In 2023 and 2024, Demis had some headway with a strategy. So he lobbied the U.K. government. He gave them the idea of the Bletchley Summit, which was the first AI gathering. China showed up to that. That then led to a follow-on summit in South

Korea. So there was some traction, also creating AI national safety institutes both in the U.K. and the U.S. and Singapore and Japan and a few other countries did this. So there was some positive movement on AI safety. It stopped when Trump was elected. I think we have to hope that when we have a new administration in this country that can be unlocked.

GILLIAN TETT: Nili...

QUESTION: Thank you for being here and for your book. Nili Gilbert. I was so fascinated by what you described about the philosophical and technological history of DeepMind, the reinforcement training, and the deep learning and how that differs from large language models. As we look into future pathways and even the implementations of the large competitors today, do you think that there could be differences in terms of implementation, future applications that continue to persist when we compare Open AI to DeepMind and others because of these philosophical and technological intellectual differences? Or at this point are they kind of competing peer-to-peer mostly for compute and the broader race?

GILLIAN TETT: And I'd argue, the question, are we going to see today's obsession with LLMs turn out to be similar to everyone chasing Betamax that then gets overtaken by VHS? Because quite seriously, because there are other forms of AI which are coming

up, you know, outside Silicon Valley in particular, but also in Silicon Valley which could end up displacing much of what people are betting on so heavily with LLMs.

SEBASTIAN MALLABY: Yes, I'm quite bullish on sort of LLM dominance in the sense that that architecture, the Transformer architecture, is a fairly multi-use architecture for making sense out of sequential data. Any sequence, whether actually that's a DNA sequence for protein folding – they use a Transformer model in the end – or whether it's a language sequence in a language model, or a sequence of video clips, any sequential data is very well understood by a Transformer model.

And large language models are one kind of representation of that, but it turns out that the motor inside robotics is going to be, you know, at least partly LLM. So I think it's true that we're moving beyond just LLMs and we're going to add new things like world models, understanding of the physical world, and that's going to make everything much better, especially for robotics. But I don't think that means throwing away the LLMs personally. That's my bet.

On Nili's question about whether we see differentiation, I think there was a view about four years ago when the labs stopped publishing their scientific work, that you would get differentiation. Hitherto, all the scientists and all the labs just published in Archive. So everybody knew everything about what was going on, and so there was no

differentiation.

And the notion was now you're going to have IP and different labs will develop different IP and they'll go in different directions. I think what stopped that from becoming true so far is that there's a lot of job hopping by the scientists and they move around. And although there are certain rules about, you're not supposed to show up at your new employer and totally spill the beans, in practice you'd give a few hints and it's quite clear what you're saying.

So I think so far, the differentiation hasn't been there. I think maybe the most likely way it comes about is you see companies like Open AI now having to make choices for financial reasons. They just can't afford to do everything. And so they closed down their video generation model Sora because it's very expensive to generate video and people won't pay you because mostly, they generate video just to put some jokey little clip on X. And so if it's just slop, nobody's paying for that. So now you have a product that costs you a huge amount to provide and of course you want to close that.

So I think different labs will perhaps make product choices that vary and they'll find different niches. And Claude is very good at coding and maybe when it comes to video generation the deep-pocketed Google DeepMind will persist and dominate that. And they've got YouTube already after all. So I think some differentiation will come in.

GILLIAN TETT: Do we have time for one last question? Okay, one last question, and very, very quick indeed. We're almost out of time.

QUESTION: Thank you. You touched upon Demis quite extensively but the computation at Anthropic, I recently had the opportunity of meeting with ____ and Bill. Do you see who is going to win when we come all the way to singularity? What kind of AI race, whether it's Gemini or Anthropic or somebody else?

SEBASTIAN MALLABY: I think Anthropic's Claude and Gemini have both had moments in the last six months when they were deemed to be the best. So the last part of 2025, Gemini was thought to be the best. Now in early 2026, the Claude model that was released in January is thought to be the best. And I think, you know, ChatGPT, don't count them out either. So there's going to be a race where at different points, different ones are leading. I honestly don't know which is going to win. I would just say that in terms of business model, Google has the deepest pockets. And so if it's up against labs like Anthropic and Open AI, which are resource-constrained relative to the amount of compute that is needed, I kind of have my bet on Google for that reason.

GILLIAN TETT: Well, thank you. I mean I'm going to; I found this absolutely fascinating. I would urge you all to read the book because it's a ripping good yarn, as they would say in an English pub. And it sort of really, it's three key points. Firstly, the power of

immigrants to drive progress and innovation, not just in the U.S. but around the world, and that's frankly something we should all be standing up and screaming to the high heavens about right now. Because this really is a story of extraordinary resourceful and tenacious immigrants beating the odds over and over again.

The second thing I find fascinating is that, as I said before, in an era where we're so focused on Silicon Valley, not all the growth and dynamism and innovation right now is in Silicon Valley. There's some really interesting stuff happening around the world as demonstrated in this story.

And the third point is that this really is a very fluid moving race. And some people are ahead at certain points both technologically and financially and then someone else sneaks up. And right now with everyone saying that they are racing to AGI, I think it would be extraordinarily foolish for anyone to bet too strongly on who is actually going to seize that crown first or even what kind of technology is going to actually win out in the end. It's really, it's one of the most fascinating and existential business and technological and people-based battles that most of us have ever seen.

So thank you so much for being with us, and over to Bob...

CHAIR ROBERT K. STEEL: Great. Well, thank you both, Sebastian and Gillian. I came

to this with high expectations which ended up being dramatically exceeded. The book is fascinating, and Gillian, you did a great job being interlocutor. And so thank you to both of you.

You can see on the back of your programs at your place that we have a good upcoming run of show for your Club in the weeks ahead. Notably, Jenny Johnson, the CEO of Franklin Templeton, is on April 21st in Florida, for those of you snowbirds that are heading south. Then if you look ahead, whether it's Lloyd Blankfein, whether it's Tony Blinken, Orlando Bravo, and also Jonathan Haidt, who I think will be fascinating, are coming up as programs here. So not to leave anybody short, but those are some real highlights. You can read all of them which are on your program. And David Solomon, I beg your pardon, so we needed to get Lloyd and David, so we kind of got both ends of the story.

Also too, let me say thank you to the Centennial members who provide support for our Club – 447 of our members have agreed to provide additional support so that we can be sure to have the financial sound footing to continue our programming. Thank you both again so much. And thanks to all of you for being here. Those online, thank you and have a wonderful afternoon. Those in the room, please enjoy your lunch. Thank you so very much.