

## Understanding AI: AlixTalks with Simon Freakley, Michael Kratsios and Amna Nawaz

With Simon Freakley, CEO, Alix Partners, Michael Kratsios, Managing Director at Scale AI, and Amna Nawaz, Co-Ancor of PBS NewsHour

**00:06**

**Amna Nawaz**

Hello, everyone, and welcome to another edition of AlixTalks, the series in which business and thought leaders gather for conversations about leading through, managing and anticipating change in a global economy driven by intensifying cycles of disruption. Today, we're going to be focused on understanding artificial intelligence and what it means for the future of business. I'm Amna Nawaz, co-anchor of the PBS NewsHour, and I am delighted to be here today in conversation with two leaders you already know very well. Simon Freakley is the Chief Executive Officer of AlixPartners. Simon, good to see you.

**00:41**

**Simon Freakley**

Great to see you, too. Thank you.

**00:43**

**Amna Nawaz**

And Michael Kratsios is Managing Director of Scale AI and formerly the fourth Chief technology officer of the United States. Michael, how are you?

**00:52**

**Michael Kratsios**

Great. It's great to be here. Thank you.

**00:54**

**Amna Nawaz**

Excellent. So, just a quick note to our audience out there. Thank you for joining us. We are going to be taking your questions in real time. So, if you want to add them into the chat, I'll try to work them into the conversation. And if there are a few others, I'll save them for the end of the conversation, and make sure to get you answers to the questions that you have.

But, gentlemen, let's jump in here, because there is a lot of conversation, a lot of fear and concern around issues of AI, I can say certainly in media and journalism, a lot of questions. So, after the release of some early platforms like ChatGPT, there was a lot of excitement, but there was also some panic. There were headlines like, should we fear AI? And is AI coming for your job? There is sort of a pervasive sense of skepticism around some of these conversations. So, Simon, let's begin with you. Are you hearing that same kind of concern and skepticism among business leaders? What's the general attitude around the AI conversation right now?

**01:52**

**Simon Freakley**

Amna, I think you captured it in your opening remarks. I think there's excitement and there's

concern in equal measure. Excitement because of the potential to use generative AI to drive a business's growth or capture more market. But on the other hand, the threat of what it might do to undermine a business, a whole industry or employment generally. So, I think in equal measure excitement and concern.

**02:16**

**Amna Nawaz**

And are those conversations sort of around excitement leading towards people to hurry towards incorporating it, or what's the willingness level when it comes to implementing AI in their businesses right now?

**02:28**

**Simon Freakley**

Well, one of the great privileges of my role at AlixPartners is I get to chat with chief executives and C-suite leaders every day, and I hear firsthand a very common conversation which is “We don't quite know whether we should be an early adopter, whether we should be a slow follower, how much of the investment that we make today is going to be relevant in two or three years' time?” So, I think that every executive is dealing with this question. I think everybody understands inevitably it's going to transform every company in every industry in every geography. But what does it mean for them?

And I think what we're seeing emerging as a theme is that most people have decided to be slower followers rather than fast adopters to see, as they watch others move in the market, what seems to get traction or what might be a cul de sac in a strategy. So, I think everybody's paying very close attention, knowing they have to have a strategy, but are not moving too quickly in case they become a hostage of a blind alley.

**03:33**

**Amna Nawaz**

Michael, what do you make of where the conversation is right now, especially some of the fear, some of the concerns around implementing it too early? Is that misplaced, in your view?

**03:44**

**Michael Kratsios**

Well, I'm a techno optimist. I work at an artificial intelligence company. I deeply believe in the power of this technology to transform businesses and to really accelerate whatever companies are trying to do – whether they are trying to get out to customers better, whether they're trying to increase sales, whether they're trying to increase efficiencies, there's so many different use cases. But one thing that we consistently hear, though, with customers is, as we come in excited with our solutions, what often comes back is this question of “How can we do the appropriate testing and evaluation of these models to ensure that what we think they're doing is what they're actually doing, and how can we measure the impact that they're having?” And I think those two questions still remain kind of as an open science, if you will.

If you listen to any of the leaders of the large language model builders, when they testify on the Hill or when they speak publicly, they will publicly acknowledge that the testing and evaluation regime that you need in place for these LLMs is still a very much open question. So for us, what we're building at Scale and what we're hearing across the industry is “How can you deploy these models,

but at the same time do it in a way that's safe and in a way that the end customer understands what they're actually doing?"

**04:52**

**Amna Nawaz**

You've raised so many important points we're going to pull apart over the course of this conversation in more detail, but I wanted to put to you this issue that Simon raised about the sort of a slower adaptation here in terms of being a slow follower versus an early adapter. One of the things I've certainly heard from tech leaders in particular is, if you're not sort of first in the space, then you risk losing the race. That seems not to be the case when it comes to implementing AI and business practices. Is that fair to say?

**05:22**

**Michael Kratsios**

To me, I think it's very industry and use case specific. So, there are certain industries that sort of have been shaken to their core with the advent of ChatGPT and the reality that their entire sort of business model is threatened because of it. Take, for example, the world of EdTech, when you used to be able to sort of go online and use a number of educational services to be tutored, those have now seen a pretty massive decline in usage in just the few months that came after ChatGPT came out. So, I think in areas like that, we're seeing some pretty quick adoption because they realize that something needs to be done in order to not be left behind. So, I think it's very use case and industry specific. I think in some places there's a lot of pressure to move ahead very quickly. In other industries, I think especially in sort of highly regulated industries, in places like financial services or healthcare, rightfully so, there's a lot of care being taken and ultimately, I think, being slower adopters.

**06:24**

**Amna Nawaz**

Simon, Michael uses the phrase shaken to their core for certain industries. I wonder if you've heard that reflected from some specific leaders and what specifically that means?

**06:34**

**Simon Freakley**

I think absolutely. Many of us, I think, realized that there was an asteroid somewhere out in space, many millions of light years away, and that in our common consciousness, that AI was growing and what we used to call machine learning was growing into something more and bigger. And then all of a sudden, we looked up and saw in the sky this planet killer that was about to hit us and the impact on the common consciousness of what generative AI, as of course demonstrated by ChatGPT hitting so dramatically, so fast, I think became something that was an extraordinary impact on all of us. What could this do to every aspect of our lives, not just our businesses? Will jobs as we understand them exist in the future? Will the dark side of AI take over? Lots of speculation, of course, about that.

So, I think maybe because ChatGPT had such an extraordinary and dramatic impact on the common consciousness with its release, I think the most adopted app ever, people suddenly realized the reality of this and that it was here and now. And so, for all of us in our personal lives, in our business lives, it's been a topic of conversation, the thought that everything will be reimaged. So, yes, I think looking at it through a business lens, I think every single business leader around the world is thinking

“What will this mean for my company? What will this mean for my industry? And how do I thoughtfully, planfully,” and as Michael says, at speed in some industries, “work out what this means for me? What does it require by way of investment? How should I be prioritizing my investment and maybe de-emphasizing other things? Is my boardroom fit for purpose in this respect? Can they ask me the type of questions to challenge me on whether I'm taking the right steps or not?” I think there are profound and multiple questions that people are dealing with, knowing that they will have to work out what this means for them. How will they capitalize on the opportunity, how will they defend themselves against the risk?

**08:38**

**Amna Nawaz**

I love this analogy of an asteroid sort off in the distance heading towards us, because I think for many of us, it did feel that way. There was a conversation about, “This is coming, this is coming” and as you say now, it is here. So, let's talk about where we are in this AI revolution and how the timeline sort of moves from here. In both of your views, Simon, I want to put to you some really striking numbers. I saw in the latest AlixPartners survey of CEOs worried, basically that their companies aren't adapting fast enough – some 75% said that they had that concern. 85% of those CEOs said they don't even know where to start. So how do you begin those conversations?

I know Michael mentioned that it is, of course, industry specific, but in talking to business leaders about those, they're concerned they're not moving fast enough, but also not knowing where to start. Where do those conversations begin right now?

**09:32**

**Simon Freakley**

Well, it's such a great question, because I think one of the challenges for business leaders around the world is that it is not just that there's one big disruption or challenge, there are multiple disruptions and challenges which are layering on top of each other, turning a rolling sea into a choppy sea. It's almost overwhelming for people and we see this in the disruption index findings. Of course, we saw, talking about statistics, we saw a 20 point jump in CEO anxiety about whether they'd be able to hang onto their jobs or not. They feel so overwhelmed by these disruptions. When we hit COVID in 2020, we thought this was the profound disruption of our time, since which we've had the tragic situation in Ukraine, the extraordinary challenging situation in the Middle East, and now, of course, generative AI as well. Different in character, but profound challenges and disruptions at every level.

So, I think the CEOs are feeling overwhelmed. They thought they were in the hot seat, but now find they're potentially in the ejector seat. The average tenure of S&P 500 CEOs now is under five years. And so, the challenge for leaders to lead through this time of profound disruption is probably the greatest it's ever been. But back to AI and generative AI, this is something that is not just a once in a lifetime transition, but probably in the history of technology, the most profound disruption that we've seen so far. But again, every one of these technological moments has both opportunity and threat in it. So, every executive is thinking, well, “Where is the opportunity for me? Where is the threat to me in this?”

One practical example is what does it mean for the future of employment. Of course, there is much speculation that generative AI will displace in a profound way employment as we understand it. My own view, and certainly talking with business leaders, our clients, is that it'll redistribute jobs rather than replace jobs. A practical example, of course, is back in 1994, nobody could have imagined that

there would be today, in 2023, over five million web designers in the world. Why not? Because there wasn't a web in 1994. And I think we'll see this with generative AI, we can't imagine how the workforce will change in nature. I personally am not concerned it will displace labor as a whole. It will simply redistribute jobs. And as a practical example, I think every business leader is working out how to redistribute jobs in their enterprise to rise to this challenge.

**12:13**

**Amna Nawaz**

Michael, talk about how you're seeing that at the implementation level as well. Those concerns are real about employment displacement. How are business leaders – you're talking to consumers, you're talking to, looking at that and what other avenues, what other lanes do you see AI, as you mentioned, we're already in this. Do you see it being implemented effectively? Revenue or supply chain? How else are you seeing it already in the works?

**12:39**

**Michael Kratsios**

Yeah, I think most of the customers that we interact with are excited about trying to understand where artificial intelligence can make the biggest impact for their business and do so in sort of the least strenuous way to get to some sort of answer. And what's sort of been very unique about what's happened with these LLMs is that the path to implementation now is materially easier than it was prior to the large language model revolution. So, if you think about a traditional AI deployment, maybe two, three years ago, it was actually a pretty complicated process. You had to find data. You had to prepare or annotate data. You had to train a particular model. You'd have to test and evaluate that model. Then you have to see where it's not working, then go collect more data and go through this cycle. And this oftentimes was a multi-month or multi-year process.

So, a lot of large enterprises sort of, I think, that attempted to do this, and you probably know this best with customers you've dealt with in the past. They realized it was an expensive endeavor that took a really long time. What large language models have allowed now in the industry is for much quicker implementation. So, these large language models are known as foundation models, and they are just that. They're a foundation that then you can build on top of with data that you have at your firm. And what's amazing about these is you can take any natural language, whether it be in PDFs and emails, in your share drive, and immediately be able to create a customized large language model for your business.

And because of that, what we're able to do and what others in this space are, you can ultimately get results for customers in a very short amount of time, in a matter of weeks rather than months or years. And I think in that sense, you can sort of show to higher level executives and people of the company that progress is happening. So to me, I think what I'm most sort of excited about is being able to sort of move through with a variety of customers and show them that you actually can do something pretty quickly now with AI. And I think that's something that you haven't really experienced in this space to date.

**14:35**

**Amna Nawaz**

Talk to me a little bit more about that speed of the process, though, because as we all know, I mean, the fine tuning of some of these platforms is where the devil is in those details, right? And we know

the models are only as good as the input. So, if you're moving quickly, is there a potential you're missing inputs that could be disastrous down the line?

**14:55**

**Michael Kratsios**

You bring up a very good point. So, I mentioned this a little bit earlier, but I can't stress that enough. The most important aspect of enterprise-grade large language model deployments is a robust testing and evaluation regime. So, there are lots of vendors in the space, lots of people who are sort of working in this. How do you fine tune a model to particular data of a company? And the most important thing is, while you can sort of throw some data at it and hope that it works out, we've all experienced moments when we've asked ChatGPT a question and it's hallucinated and given us an answer that we never imagined possible or don't understand where it came from. And in most enterprise use cases, that is pretty unacceptable most of the time.

So there has to be a lot of work on this test evaluation piece. That's something we think about. That's something that the industry is trying to perfect the moment, and we may get to this in a bit, but you even saw in President Biden's executive order last Monday, there was a lot of talk around building standardized methodologies for testing and evaluation that the government can work on. So, I think everyone recognizes that no matter where you deploy these things, having a standardized way to think about test and eval is front and center important.

**16:01**

**Amna Nawaz**

And I want to get to the regulation question in just a moment because I think there's a whole different conversation around that. But Simon, I want to get your take on what Michael just said, especially about some of the concerns, potential harms is, of course, a willingness and a desire to move quickly to do what's best for your business as a CEO. But how are they thinking about the potential harms, the fact that employment, as you said, redistribution could be happening? Where are those future jobs? Who are they looking to hire and bring on board to make sure that those potential harms and disastrous consequences from not having the right inputs don't happen?

**16:36**

**Simon Freakley**

Right. Well, what I found just in my many conversations with client CEOs is that the very best leaders are just really getting to the nub of the key questions for their business and their industry. What are the key questions they need to answer to be able to make sure that their business remains not just relevant in the future, but a high performing business in the future? What is the nub of the key questions that they need to be in a position to answer to work out what they should do in terms of next steps as they develop this technology within their business? And of course, most chief executives, because of the tenure where they're at in their careers to be chief executive, haven't got a digital background. Some have, but many don't. They're not digital natives. I mean, they work very hard indeed to understand enough to ask relevant questions, but the very best leaders put people in the room with them that really are digital natives and understand some of the key questions to challenge the leaders with.

So, assembling the right team to be able to get to the nub of these questions is key. And then working out from an investment profile the type of investments that are not just desirable but are also affordable. I personally think that every corporation, major corporation, will have its own large

language model, probably several large language models, to do different things, that those models are unlikely to be developed from scratch by them. They'll probably take a proprietary foundation model and then adjust it, make it bespoke to their enterprise, train that model on their own data so they can synthesize both information from public source and information behind their firewall. And that's what I find most leaders are doing at the moment, working out how to navigate to that position, how to do it safely with several choices, which would be the best choice for them, but putting people in the room with them, that will likely enable them to get to the right answer, because probably their own experiences are not going to be enough of an insight as they make these critical choices.

**18:42**

**Amna Nawaz**

Simon, I'm just curious, most of the business leaders you've talked to, did they experiment with ChatGPT and check it out when it was released?

**18:51**

**Simon Freakley**

Mostly, yes. I mean, as I chat with people, it's quite interesting to hear some of the stories that they will use to illustrate the point, some of which, of course, are very much business questions. One I spoke to recently used ChatGPT to make up bedtime stories for their children. Somebody else was looking to sail quite a long distance on a boat and asked ChatGPT to tell them the best way to do it, which was remarkably accurate, as it turned out, when they then compared it to a maritime service. And so, people are experimenting, I think, to learn how to use it. And I think that as they experiment, they realize that it's not like a Google search, that actually you can ask some really quite sophisticated relative questions, not just absolute questions. And I think that this interaction that people are having are helping them develop a frame of reference, but understanding they'll never be an expert, and they need to bring the experts into the room to help them understand the real questions they need to be asking.

**19:50**

**Amna Nawaz**

So, Michael, for those leaders you talked about fine tuning, rather Simon mentioned, starting with some of the open source models and then fine tuning them to their own businesses and needs. How does a company know they've reached a level of confidence that it's ready to be rolled out, it's ready to be implemented? How should they be thinking about that?

**20:11**

**Michael Kratsios**

Yeah, I think it goes back to this test eval sort of question, and I think what we typically advise and try to work with customers out of the gate is on sort of non-mission critical solutions. So, the first implementation doesn't need to be something that is sort of like do or die for your business. You have to be in an environment where you can imagine that the end user, the LLM, probably the company will be fine if some hallucinations come through and those can be evaluated and monitored and recorded, and the model can just improve over time. So, I think for us, there's a battery of tests you can do that allow you to kind of check the performance of the model. You obviously can have it running for a certain amount of time. And what we found is that the greatest experts are actually people in the company themselves, in the customer themselves, that are using this model every day, and they can provide extraordinarily valuable feedback on where the model is

working, where it's not working, what areas, where it's not performing, where it is performing. And then from there you can go out and generate more data to help train the model in the right direction. But I think the top level, 30,000 foot view is trying to find use cases that are important for the business but aren't going to create a huge headache if they go wrong. So, sort of basic, you wouldn't want to create a large language model for an outside customer-facing solution out of the gate. You probably want something internal for your own teams to be using first and then slowly grow into something more external.

**21:42**

**Amna Nawaz**

Can I ask you something based purely on a consumer standpoint? We hear this a lot from people with consumer facing products in particular, is that this kind of development model sort of relies on things going wrong first before they get fixed, right? If the inputs haven't already been implemented into the model. And I wonder how we should think about that outside of business leaders, because as we keep mentioning, the platforms are only as good as the inputs. So, it feels a little bit like kind of building the plane as you're flying it. Is that fair?

**22:13**

**Michael Kratsios**

A little bit, yeah. I think both when it comes to regulation and even when you're trying to implement these in a corporate setting, it's all about taking a risk-based approach to giving the green light for deployment and there's very sort of simple, basic, easy things that you can do, that the world's not going to end, everything's going to be fine if it doesn't work out. And there are other ones that aren't. So you, as a leader in a corporate setting, have to be very thoughtful around trying to create a risk framework yourself on what are the use cases and how do they fit within sort of the importance of the business. And then from there, you can figure out how much testing and evaluation preparation you need to do before you hit the deploy button.

**22:52**

**Amna Nawaz**

Related to that, there is this issue of internal biases, how they get implemented unintentionally, unwillingly, sometimes into the platforms and the work that we put out. There's actually an audience question that just came in that helps me make this pivot. Someone watching, thank you for your question, has written, "How are leaders looking to mitigate any biases that can arise in the LLMs?" Simon, let's just start with this. Is this part of the conversation? Is this something that leaders are cognizant of among the CEOs you talk to, that it can fuel efficiency, AI can fuel efficiency, it can boost competitive, solve this. But the consideration of impact and potential biased impact on consumers, are they talking about that?

**23:37**

**Simon Freakley**

They absolutely are and, of course, it's not as if they've just started thinking about it with the advent of generative AI. The whole question of is there systemic bias in all sorts of processes, from the obvious one of recruitment, for instance, but all the way through the different business processes, how does one spot bias? How does one adjust for bias? And this has been a very live and important discussion that's now been going on for quite some years. So, I think the discussion that we're now having on the back of how large language models are programmed to ensure that they don't display or inherit the bias that's there in the data, but actually further cleanse the bias out of the system, is a



very real discussion, and I think a very important discussion. I must say, Michael, might have a different view, I haven't heard of anybody making a breakthrough in terms of how to do that in the programming of LLMs or the sanitizing of the data on the back of which these programs learn. But what I am really pleased about is that the whole discussion, as you say, about how do we use this as an opportunity to further take bias out of the system, is a very live debate and I think a very important debate.

**24:47**

**Michael Kratsios**

Yeah, what we tend to see is that if you are seeing or evaluating bias at the end of the road, at the output of the model, you probably already have lost. The way that you tend to think about this is you need to be evaluating and thinking about these questions at the beginning of the process, on the initial data that you're using to train the model itself. So, we have spent, Scale as a company has spent many of the years, even pre LLMs, thinking about this problem and the best way we have found to solve it is you have to run your testing and your evaluation on the initial data sets before the model is even trained. And at that point you can have a good sense of what does the data look like? Is it skewing in a certain direction that doesn't seem right for whatever reason? And then from there, once the data is in a position where you think it accurately reflects the direction that you should be going for your particular use case, then you can start training the model. So for us, it's data first, always, and if you're worrying about it ten steps later, it's a bit of a mistake. You may be too late.

**25:54**

**Simon Freakley**

Just in terms of a pre digital example, and I'm sure you're aware of the example, but classical orchestras, classical music orchestras were just appallingly undiverse. And so, people started to rightly ask the question in 70s, and the 80s, and then the 90s, why is it that the best orchestras of the world are almost entirely populated by middle class white men? Where's the diversity in classical music? Surely there are brilliant musicians, female musicians and people of color. And so, the main, most prestigious classical music orchestras in the world started to do these blind auditions, whereby they would have, when they were auditioning for chairs in the orchestra, they'd have people come onto the stage behind a screen, so they couldn't see their ethnicity, they couldn't see their gender. And then people realized that unconsciously, people were picking up from the click of the heel as they walked across the stage behind the screen, whether it was a female heel or a male heel. And so even though it wasn't conscious that actually the unconscious bias was still taking shape. And so, then they further enhanced this by having a carpet where the person could walk to the spot, so you couldn't either hear the click of the heel or see the person. And surprise, surprise, here we are in the 2020s and we have much more diverse orchestras today than we had even 20 years ago. And so, the very deliberate act of seeking how bias can creep into decision making and then doing repetitive steps to take it out, whilst in a nondigital world the classical music story illustrates it, in the digital world, there are just different techniques. I think we have to be as deliberate in our attempts to keep the bias out of the data on which these models learn as we do about the output of the choices. And I think this is, what I'm thrilled about personally, is that we are very focused on this as a Western society to make sure, to the extent possible, we use this as a way of averaging up our performance to take out bias from decision making rather than just repeat our errors of the past.

**28:09**

**Amna Nawaz**

There's a number of audience questions coming in I want to get to as well, but we talked about regulation a bit earlier, and I'd love to get both of your takes on this, because, of course, here in the US President Biden has just made history, right, signing the most expansive regulatory attempt yet with a sweeping artificial intelligence order at the end of October. He's essentially invoking broad emergency powers to harness the potential and to tackle what he says are the risks of what he called "the most consequential technology of our time". Michael, let's just start with you and your reaction to that executive order. And in the US in particular. Is that kind of regulation necessary at this stage of development, and what do you think will be its impact?

**28:54**

**Michael Kratsios**

Yeah, it was an interesting order. I think the key question that the Biden administration tried to solve for and out of the gate, most prominently was what do you do about existential risk posed by these large language models? So, the question of could a bad actor potentially use one of these models to ask a simple question like, "How can I create a biological weapon using materials I can purchase on Amazon?" Things like that, where you could essentially, they are low probability, but very high impact risks. And the question is, can the government, and should the government be in a position to essentially work with these companies to ensure that type of risk is minimized?

The Biden administration took the strong position that they should be involved, and all the major large language model providers are now required to be sharing their testing data with the federal government. The second piece of executive order thinks a lot about how you can actually start building a regulatory regime around these large language model use cases. And it did not call for a new AI agency of any sort, but it did call for a number of agencies to start beginning the hard work of thinking through what a regulatory regime would look like. Most prominently, the standards agency of the federal government, called NIST, has been tasked with creating a set of standards, if you will, for testing and evaluation of these models. These, I think, could be very valuable for the larger industry, where we can all be singing from the same song sheet. When it comes to before you deploy something, what do you end up testing? But I think there is a bit of a question in some parts of the industry whether or not this is a little bit too early. If you can think a little bit about the general take that the US and Europe has had. Europe has always taken sort of this precautionary principle, this idea where you have to think a little bit about the harms and maybe create regulations to try to minimize those before those harms have been materialized or realized.

And in some sense, I think many people are viewing this executive order maybe a little leaning in that direction where this is great technology, we know some things may go wrong, like let's start thinking about those now and regulate it. So, there's a big push and pull. But I think the big fact is that these large LM providers are now required to disclose certain things to the government, which is a pretty big action forward.

**31:06**

**Amna Nawaz**

Is it fair to say the US is far ahead out in front of everybody else when it comes to this development?

**31:11**

**Michael Kratsios**

At least on the regulatory front, the European Union has been working on an AI act for many years now. So, it almost felt like the US was a bit trying to sort of catch up with it. If you're in a race to who

can regulate first, I guess the US was trying to catch up. I think many people would say maybe it was probably better that the US set out a little bit. But there was a large summit in London, outside of London last week, where leaders from most of the large countries around the world came together with the big large language model developers, and there's a sort of global consensus that thinking through some of these risks is actually really important. So, the US is kind of charging ahead. Now the question is how is that going to impact individual corporate use cases? I think that's still a little way out and to be seen

**32:00**

**Amna Nawaz**

Simon, how are CEOs viewing what's unfolding here in America, both on the development side and also to see how this regulation plays out?

**32:06**

**Simon Freakley**

I think everybody is trying to assimilate as quickly as possible this developing landscape from a human point of view, appropriately what regulation is doing to protect society as well as business, from some sort of dystrophic outcomes of what might happen, but also how to understand the regulatory environment so one can operate within it. And so, I think that people are reading as much as they can as fast as they can to orientate themselves, but understanding that there are specialists like Michael and others that are at the forefront of this, and they are hungry for insight and advice, how they should best steward their companies through what is a fast developing landscape.

**32:50**

**Michael Kratsios**

I think one thing that's sort of, I think, positive for the larger corporate community is that the general trend in the US on AI regulation is that it should remain industry and use case specific. So, if you are a medical diagnostics company and you want to create an AI powered medical diagnostic, it will still be approved by the FDA. There may be new tests that need to be done because it's AI powered, but there won't be a new agency you have to go to. At least that's generally where Washington is moving. So, for folks who every day interact with the FDA or interact with the Department of Transportation, your regulators will likely stay the same. The question is now, how will AI be incorporated into the work that they're already doing?

**33:32**

**Amna Nawaz**

I want to get to some of the wonderful audience questions that are coming in right now. And thank you to all of you out there, keep sending those in. I'll work them in now. Simon, let's put this one to you. One of our audience members is writing "Which types of use cases do you see where businesses will purchase external AI services, and where do you see them developing them internally?"

**33:53**

**Simon Freakley**

I imagine a situation where every major corporation is going to have its own large language model, or maybe several. I mean, I spoke to a chief executive last week, for instance, an organization employing 120,000 people. Now, they were in the tech space. They are currently working on 20 large language models of their own. Now, I think that some of those have been foundational models

they've taken from others and are building and training on their own data, some of which they've developed from scratch. And so, I think that depending on, as Michael said earlier, the sector and the urgency, I think people will be looking at either buying or building, or a combination of the two, sooner rather than later. I'll give you another example – I spoke to one of our private equity clients recently, one of the very large private equity houses, who were doing deep diligence on a number of major deals. They'd licensed from Microsoft a ChatGPT engine that could work behind a firewall, so they could compare, they could use the engine within the data room on these diligences, as well as synthesize the open source information as well. They'd done a number of deep diligence exercises on deals that they'd already done, but ran this engine on the old data rooms to see what else the engine would highlight that their own diligence hadn't. And in every case, according to this senior partner, the ChatGPT engine under private license in the data room, had come up with at least three material insights that none of their diligences had.

So, if you're a private equity house to this example, you want to have a model that works in a behind the firewall environment. You really don't want to be building one yourself, you'll take a proprietary one, maybe ChatGPT, or one of the others. And so, I think that for those that need to use these engines urgently or see a real premium in moving to action quickly, then licensing an existing engine to be used behind a firewall will be the priority, albeit in time they might develop their own as well. For some other specialist businesses, I can imagine they will write their own.

So, one of the great opportunities, I think, for generative AI is in the development of medicines, the development of treatments, and the speed at which generative AI can seemingly iterate different use cases for medicines seems to be an extraordinary benefit for humankind, let alone for the drugs industry. And so, I think that people will be using proprietary models for that. So, I think it's somewhat use case dependent, but I think buy, build, and both is going to be on everybody's agenda.

**36:40**

**Amna Nawaz**

Michael, anything you'd add to that in response to the audience question?

**36:43**

**Michael Kratsios**

No, I think that's right. I mean, when I fast forward a year or two years, most large enterprises, I think, are going to have many tens of LLMs that are fine tuned to the particular use cases of the company. And the base model that you use for those LLMs is probably use case specific. So, there's some models that are essentially open source and free, and you can use those for certain use cases, and they may be good. There are other more high stakes use cases where you want to use a more expensive, cutting edge frontier model. So, there's a lot of actually interesting work being done on trying to help companies identify which foundation model they want to use for their particular use case.

**37:22**

**Amna Nawaz**

There's a number of questions coming in on the issues of oversight and company level regulation, if I want to put those to you as well, because I think there's a lot of interest in this right now, as it is uncharted territory. But Michael, let's start with you on this. One of the audience members is asking,

“Does AlixPartners see companies looking to have an AI ethics officer like IBM has done?” Is that something you see companies putting into place?

**37:49**

**Michael Kratsios**

As I have talked to a number of our customers, there are a variety of approaches that they're taking in order to evaluate these AI use cases. So, IBM and I think Microsoft and others have taken sort of this almost council-type policy approach where there's a group of senior leaders, some of them are deeply technical, some are more sort of like legal in the ethics domain, and some are in the business domain, and all them together can make important decisions about where and when to utilize these models. For very large companies like that, they're deploying lots of models at very large scale, so that type of structure probably works. I think, for small organizations or for ones that are just dabbling with AI, I think they can potentially grow into that. And back to the ethics officer, I think within the legal domain of a lot of these companies, we're seeing those roles pop up. At least we are. I don't know what you're seeing, Simon?

**38:47**

**Simon Freakley**

I completely agree with you. And I think ultimately, absolutely every major corporation will have to have somebody fulfilling that responsibility. And of course, there are different standards, and it's not just about meeting the lowest standards. So, there are legal and regulatory requirements. But I think your question goes beyond that as to what are the ethics that should govern how data is amassed and synthesized? How is that ethical test applied to the values of the corporation? Of course, as we know, for all of our stakeholders, particularly our employees, but also our customers, our lenders, our investors, the values, the purpose of an organization is key. And so, being able to demonstrate in this particular dimension that the ethics of an organization, by reference to its values, make sure that data is both aggregated and then used in an ethical way is going to be critical. So, in the same way that we see in certain Western economies, chief money laundering officers or chief diversity officers or chief information security officers, I think we're absolutely going to see people in this role, and I think it'll be a very important role.

**39:55**

**Amna Nawaz**

Simon, that's fascinating to hear you say that. You think every big company will have to have someone in this role to oversee this, is that right?

**40:02**

**Simon Freakley**

I absolutely do. I think it'll be not just to meet regulatory and legal requirements, which is the threshold requirement, but to be able to demonstrate to stakeholders this is something that the corporation takes seriously and is focused on, and the appropriate reporting structure is in place so that somebody fulfilling that responsibility can escalate to certainly the chief executive and probably the board very readily if they fear that information is being aggregated, synthesized, or used in a way that isn't consistent with the firm's values.

**40:36**

**Amna Nawaz**

So, this actually gets to another audience question coming in that I'll put first to you, Simon, and

then Michael, I'd love to have you weigh in. Simon, you talk about every company having someone like this in their hierarchy who's sort of overseeing all of this. One of the audience members is asking, "On the issue of training against bias, who will be the keeper of the rules?" So, how do you sort of manage that within a company, how are CEOs thinking about that, Simon?

**41:01**

**Simon Freakley**

Well, inevitably, what happens, of course, in these situations is that these types of things tend to aggregate under the general counsel, because when there isn't an obvious answer to something like this, even if somebody's given particular responsibility, they report to the general counsel, because the general counsel is the person that looks across the landscape of risk and looks across the landscape of liability. And so, often that is what happens. Now, that isn't to say that the chief information security officer, somebody doing a very important job, particularly in this world where we're surrounded by cyber threats, isn't reporting up with a very different mandate, but for a similar purpose to make sure that the enterprise is kept safe and threats are mitigated.

But in the ethical officer, if we call them that as a shorthand, how one can demonstrate to various stakeholders that a corporation is handling data in an ethical way. I think that may sit under a general counsel conveniently, or it may be something that goes through, for instance, the head of internal audit. The head of internal audit, classically, has a direct reporting line to the chairman or the senior independent director on the board, so that person can report independently of the existing management structure. I can imagine honestly, for the want of a better title, the chief ethics officer, information ethics officer in this sense, having a similar reporting line. But I do think that fairly quickly, every corporation will need to have someone doing that function so that stakeholders can be confident in how information is being used by the corporation.

**42:40**

**Michael Kratsios**

Totally. The only thing I would add was, wherever this ends up being, I think there's certainly roles within a CIO office or CTO's Office of people who actually can do the technically rigorous analysis that then gets fed up to the particular person responsible. And I think when you start scoping out sort of, what are the technical expertises you need within your sort of CIO team or CTO team to execute on these AI projects, I think you need people who are excited about the implementation and the execution and building the model and deploying it, but you ultimately also need to think very carefully about who is doing the testing and evaluation – are those people separate and distinct from the people that are just trying to build the model – and creating even that kind of separation within your organization, so you have some sort of checks and balances on the tech side.

**43:24**

**Simon Freakley**

That's a great call out, actually. If you think about, back to my head of internal audit analogy, the head of internal audit will have a reporting line to the CFO, but actually have a direct dual reporting line to the board. And I can see a similar model taking shape. So, maybe the person does have a reporting line to the chief technology officer or CIO, but a separate reporting line elsewhere to ensure that everybody can be confident that somehow a message isn't being lost through the standard reporting structure of the enterprise.

**43:52**

**Amna Nawaz**

So, Michael, only because you mentioned it, I need to ask, do you foresee in the future at a federal level or maybe even a state level, I don't know how you manage this, something like an AI agency, because it becomes so pervasive and so intertwined in every part of our lives?

**44:08**

**Michael Kratsios**

My general sense is – I think this came up in the news right around when the first congressional hearings were happening earlier this year – but over time, what I've observed in Washington is they've sort of cooled on that particular approach. And the reality is that at the end of the day, most regulators have been working on their particular industry for quite a number of years, and you're never going to have that deep level of expertise at one centralized agency. So, if suddenly AI powered medical diagnostics had to be sort of approved by an AI agency, the speed and the quality of that type of effort would be very different than the FDA. So, what we're seeing is, I think there's been general sort of bipartisan consensus around maintaining kind of this industry specific use case approach.

The challenge about all this, and I think the executive order touches a little bit on this with its call for AI skills being brought into government, is that these agencies have never, in many cases, dealt with AI power technologies. So, the expertise that you need in order for them to make very smart regulatory decisions as technology develops is still not quite there. I think there's a lot of learnings yet to be had, but I foresee them to sort of keep the status quo with the way that agencies are structured at the moment.

**45:19**

**Amna Nawaz**

A couple more big picture questions I'd like to put to you both, and then I see more audience questions coming in – I promise you, I will reserve time to get to those because they are very specific and fascinating. But we're focusing a lot on business and business practices, I'd love to put to you a lot of the concerns I hear here in Washington, in particular among government leaders and how they're looking at it. National security officials in particular, talk about their concerns. I spoke to one senior administration official who once told me nothing keeps him up at night like AI does because it's just unknown, and there are so many concerns, especially when it comes to things like election security and infrastructure security. More broadly, I think it's fair to say, as part of the conversation around rising tensions with China in particular.

So, Simon, how do other business leaders, CEOs you talk to, how do they look at those conversations and those bigger concerns? Do they see those concerns as valid and are they feeding their thought processes and processes right now?

**46:19**

**Simon Freakley**

So, I think at a human level, we all worry about that. Of course, I don't know – many of us don't know – I have to assume that AI has been used profoundly in issues of national defence for quite some time. So, it's not as if this is a brand new development that now is going to recreate how we think about the protection of our countries, whether generative AI, of course, will create another level of threat, but also defence – I'll leave that to the experts. But I think for the business leaders that I speak to on a regular basis, of course they're worried about those issues as a citizen of the US

or a citizen of the world, but also just worried about how comprehensively it could upend their enterprise.

And so, for instance, I was talking to the chief executive of a large retail chain recently, who said that he'd just decided he didn't know, but he'd just decided that within the next five years, he needed to assume that he didn't need 30% of the jobs that were currently being done in his enterprise that he did today. And he was going to use that as an assumption, not because it was going to be precisely right, but directionally it was going to be right. And how could he retrain people to do some of the higher value jobs? How could you think about reconfiguring his business around that assumption? And he is now very much doing it, looking at the data sets that they have as a business to think about, what are the unmined pools of value in that data that by exploring that unmined pool of data with generative AI, they can create new ways of imagining revenue. Of course, every chief executive is serving at the altar of growth, because revenue growth is the most direct proxy of equity growth. And so, how to use the data they have the new capabilities that they will have access to through these large language models to create streams of revenue growth they don't have today, whilst at the same time making some assumptions to how their cost will be configured, particularly around headcount. So, I think, honestly, this is something that keeps every chief executive awake most nights.

**48:24**

**Amna Nawaz**

Michael, how do you look at this in particular. In the conversations in Washington, are the concerns they're raising valid from your viewpoint or are they largely just tied up in political tensions, right now?

**48:36**

**Michael Kratsios**

I think what I have seen is that there's been these stories or these vignettes about the way that large language models could be used to sort of create catastrophic results or pose dramatic national security risk. They get a lot of clicks, people love to talk about them on the Hill. My general perspective is these are sort of low probability, high impact events. And if you're a government and you're responsible for taking care of the country and ensuring everyone is safe and secure, yes, absolutely, there should be people working on this and figuring out what's the best way to minimize the risk, that bad actors don't use these technologies to do horrible things.

But one thing that I will say though, is you can't allow the entire conversation to be dominated by that conversation. The reality is that there are immense benefits that this technology can bring for the American people and the way of life of everyday Americans. And if we're too caught up in dealing with these low probability events, we're not focusing the energy on the way that we can potentially safely deploy this technology in a wide variety of industries. So, to me, I think there's an important balance that needs to be struck and hopefully that can continue, and we don't get too caught up in this sort of almost fear mongering, extreme risk conversation.

**49:53**

**Amna Nawaz**

Let me make sure I get to some of these audience questions that are coming in, and keep them coming in – we've got a few more minutes. I will try to work them all in. One person is asking specifically about the impact on the consulting sector. So, Simon, let me put this to you – “In light of



the increasing integration of AI into various industries, I'm curious about its overarching impact on the consulting sector. Could you share insights on how AI is reshaping consulting services, particularly in terms of operational efficiency, client service quality, and innovation in business models, and last but not least, pricing and profits.”

**50:27**

**Simon Freakley**

I think it's going to profoundly affect the consulting industry and maybe even more broadly the professional services industry in a number of ways. I mean, first of all of course, most professional services firms, including most consulting firms, are pyramids, and so the majority of the people that are junior level. And that's how the business model of these consultancies work. Very hardworking junior consultants, recording lots of hours. What's going to happen is that AI has already started to displace much of the analysis that's done by the traditional bottom half of the pyramid. And so, most consultancies will turn into a diamond shape rather than a pyramid. But the one thing that won't change is the value of wisdom and judgment, which of course is what clients look for in terms of what the analysis actually tells us, what their best strategy should be.

So, I think that the consulting industry will absolutely continue to thrive, because advice, and judgment, and wisdom is still prized in its many different forms. How the data is aggregated into analysis and then refined into what the available strategies are and how one advises on the preferred strategy, and also, particularly of course AlixPartners that specializes in the implementation of strategies I think will still be alive and well. What I do think will happen though, is that the way in which consulting services or accounting services or legal services end up being priced will be less about per diem rate, paying by the hour or the day, but more around the impact of the solutions and the implementation of those solutions. So, do I think the consulting industry, the legal industry, the accounting industry will still be necessary and thrive? Yes, I do. Do I think the way in which the work is done, the analysis and formulation work is done, to enable the senior practitioners to give their advice? Absolutely, it will change. So, I think there's going to be an awful lot of transition going on in the next few years.

**52:28**

**Amna Nawaz**

Michael, there's a good question about the development of large language models for you. Here someone writes, “Do you believe the most cutting edge LLMs will continue to increase in parameter size and training time and resources, or will they start to level off or even decrease in size while improving performance?”

**52:46**

**Michael Kratsios**

Well, I think we're going to actually see both. So, I think from most of the large language model players you're seeing, they continue to want to increase the size and scale of their models. And the next training runs that you're going to be seeing the top developers doing over the next year are going to be increasing in size. Like, right, an extreme example is the Falcon model that was created by the UAE – that's arguably the best or the second best performing open source model in the world, that's an 180 billion parameter model. That being said, I think there's a real realization that a lot of these use cases do not require these massive, massive models that can often be expensive to run.

So, there's been a lot of attempts to sort of create much smaller, tighter, less big models that you can essentially deploy much more easily and even deploy at the edge. So, I think the big model players are going to be trying to continue to build the greatest model the world has ever seen, and that trajectory is almost gone asymptotal for the next couple of years, and then there's going to be a lot of engineering being done on how you can make these smaller models as well.

**53:50**

**Amna Nawaz**

A number of questions coming in now. I want to make sure I get people answers to the questions they have. So, Simon, let me put this one to you about what companies should be doing right now. Someone writes, "What should companies do now to help with AI adoption across their users besides training on prompt engineering?" Maybe Simon and Michael, you can both answer that.

**54:09**

**Simon Freakley**

Well, I go back to something that we chatted about earlier, getting the right people in the room to ask what the questions should be. I'm not a digital native. I have colleagues, of course, that are. Angela Zutavern, who runs our AI practice, just somebody that really understands profoundly what the right questions are in a given environment and makes sure that they are in the room testing the management team to ask themselves a difficult question. That's the very first thing that people should do.

**54:32**

**Michael Kratsios**

Yeah, and to me I think the short answer is every business problem doesn't require an AI solution. And I think that's the challenge that we always see, where AI is sort of this panacea for whatever the problem is. I think a lot of work can be done on working with very smart people like Simon and others to sort of think through what are the types of problems that lend themselves well to AI solutions and make sure that your teams aren't running and chasing things that you can't actually change a lot with AI.

**54:59**

**Amna Nawaz**

Simon, let me get back to you about issues of people within the company, who's in the room and where the responsibility, accountability lies. Someone's asking, "Chief legal and compliance officers in financial services firms are very concerned about the adoption of AI tools as part of online meetings for transcription, summarization, notetaking, et cetera, because of the regulatory requirements around record keeping. So, although users are very eager to use them, what have you seen, what's your experience with this situation and what are your highly regulated clients doing?"

**55:34**

**Simon Freakley**

It's such a great question because, of course, and financial services is the right example, because the financial services companies of which have been some celebrated examples, of course, recently, some of whom have been under monitorships from the Department of Justice and others, as they try to work through some of the issues they've had, the risks in adopting this new technology too soon, without the appropriate checks and balances, the risks are completely out of kilter with the

rewards. And so, as they try to make sure they don't end up not honoring a sanction requirement, or not allowing bad actors to move money through their network, or not applying the established processes, they want to make sure that they don't fall foul of those tests, cross those hurdles. And so they're likely, I think, to hold off using generative AI, certainly open source, generative AI capabilities, until all of the protocols are in place. So, I think you'll see financial services being a slow adopter in this respect. I think once the large language models can operate behind firewalls, and so there's no sharing or potential of sharing of information open source, then I think it becomes easier. So, to the specific question, I would say that general councils of financial services businesses are actually displaying good judgment by taking it slowly, because getting it wrong could result in a calamitous fine or worse to the company.

**57:01**

**Amna Nawaz**

Michael, anything you'd add to that?

**57:03**

**Michael Kratsios**

I think that's right. I think we talked a little bit about sort of taking a risk-based approach to selecting your use cases, and I think this applies directly to this issue.

**57:11**

**Amna Nawaz**

So, we've only got a couple minutes left, but I'd love if each of you would kind of leave us with a specific example. As we know, this is not a revolution that's coming. It's one we are in and very much living through. So, in terms of the customers you work with, business leaders you're talking to, could you give us one specific example of where AI is being adopted and implemented successfully, and what benefit it's bringing to that corporation right now? Simon, would you like to start?

**57:37**

**Simon Freakley**

Let me give you a real example that happened in the last couple of months. One of our clients is a manufacturer of quite a sophisticated product requiring many electronic components, and asked us whether we could suggest a way to make this product in a more cost beneficial way. We have these tear down labs at our firm where we will take things right down, literally, to their basic molecules, to understand composition and then rebuild them. And so, in this particular case, we took the product to a tear down lab. Six consultants spent two weeks tearing this product apart, right down to its basic structures and components, and then proposed a solution where a product at least as good from a quality point of view, could be produced at a 20% discount to its existing cost.

Before we sent that response to our client, we then asked ChatGPT the question, and this is a well known public product, so it wasn't as if we were putting anything on the open source network that was at all proprietary. We asked the same question. Within three minutes, ChatGPT came back with answer that was 98% as good. It wasn't completely perfect, but it was 98% as good. And so, you start to think, well, actually, if you'd started off with the ChatGPT answer and then searched for the 2% or more, that could be improved, it would make for a very different consulting experience. So, that's a real example.

**59:03**

**Michael Kratsios**

A real example from our end. One of our customers that we love working with is an EdTech company called Chegg, and they had a bit of an issue when ChatGPT came out. It became very clear that a lot of students and other users that used to be using Chegg to be doing sort of training or educational sort of services had moved away and just started asking these questions to ChatGPT. So, the question was, what should Chegg do in that moment? And the reality is that ChatGPT is good, but it's not like that good. If you want it to have really, really good answers for sort of complex education related questions in domains like physics, or chemistry, or math, it'll be okay, but it's not going to give you the best answer.

So, Chegg had this moment where it could essentially create fine-tuned large language models for individual educational domains, so they could create a model for chemistry that will always better performant than any chemistry question you ask to ChatGPT, because it'll be trained on a large corpus of chemistry specific data. So, we've worked with Chegg to create a wide number of specific educational LLMs, and this is an opportunity where you can see that the power of fine tuning allows you to take sort of this proprietary data that Chegg had around these particular educational domains and deliver a product to their end customers that is better than anything else out there.

**01:00:23****Amna Nawaz**

Gentlemen, I thank you both so much for your time. I think it's fair to say we could easily go another hour. The brave new world is here, and we are in it. Simon Freakley and Michael Kratsios, thank you so much for joining me today.

**01:00:35****Michael Kratsios**

Thank you.

**01:00:35****Simon Freakley**

Thank you very much. Thank you, Michael.

**01:00:37****Michael Kratsios**

Thank you.

**01:00:37****Amna Nawaz**

Thanks to everyone out there for joining this edition of AlixTalks. See you next time.