

<u>Beyond Unprecedented: The Post-Pandemic Economy</u> <u>Season 4, Episode 1:</u> <u>"Business and Bots: Anticipating the Future of Al"</u>

[00:00:04] **Kareem Yusuf:** Al's superpower is decision support. We should be looking at making it clear when is Al being used. Transparency builds trust and raises the right level of awareness that guides—or should guide—how people operate.

[00:00:21] **[Music and media clips of journalists]:** The coronavirus pandemic has tanked the global economy with unprecedented speed. The steepness of the decline here is unprecedented. This is a crisis that is unprecedented. It is unprecedented, and we just don't know.

[00:00:35] **Eric Talley:** This is *Beyond Unprecedented: The Post Pandemic Economy*, a limited series podcast from Columbia Law School and the Ira M. Millstein Center for Global Markets and Corporate Ownership. I'm <u>Eric Talley</u>, Sulzbacher Professor at Columbia Law School and co-director of the Millstein Center.

[00:00:52] **Dorothy Lund:** And I'm <u>Dorothy Lund</u>, the Columbia 1982 Alumna Professor of Law at Columbia Law School, and also a co-director of the Millstein Center.

[00:01:06] **Talley:** Today, we'll explore advances in artificial intelligence and the emergence of generative AI. We'll discuss the technology's implications for the economy and impacts on workers, regulatory ramifications, and what lies ahead. So, Dorothy, we have daily lives as well. How has AI, you know, dramatically changed your daily life as a law professor and as a scholar?

[00:01:28] **Lund:** Well, thus far, I haven't seen huge changes, which is good because I'm still hoping we have a job in 10 years. I'm trying to use it more. I've been listening to different podcasts and reading books that suggest different ways that you can use it, and so I'm hoping this podcast will be the inspiration I need to really, you know, supercharge my work product with it, with AI. But I know you've had some interesting intersection points with AI and your exam writing.

[00:01:56] **Talley:** Yeah, it's been a big issue. So, for some of the classes that I teach, because I typically will give what are—what used to be called take-home exams or remote exams. And there's always a question about whether someone, even though, you know, they've signed an honor code, whether there is access to generative AI,

foundation models, and so forth, in helping to prepare an answer. And so rather than try to set up a bunch of protocols to resist it, I decided just to give in to it. And I've now, for two semesters now, been giving take-home exams in which at the end of the exam prompt, I just give the students ChatGPT4's response to my prompt, and what I tell them to do is just do better than that. And thus far it's worked out OK. But beyond that, in my classes, all these issues about how do we regulate this stuff, not just in the exam format, but in terms of intellectual property, in terms of competition law, in terms of privacy, it has become kind of like the topic in law schools, even though I usually teach in business law, I'm guessing you've seen the same thing.

[00:02:55] **Lund:** Yeah, absolutely. And I think I think the teaching idea is so clever because I think something we're going to have to start doing in the coming years is not only getting our own work and research up to speed using AI, but teaching our students how to do the same. And we actually had a really cool demonstration a few weeks ago, at one of our legal aid clinics that is starting to roll out AI chatbots to supplement some of the legal aid services that they're providing. So, you know, I do think this is going to be something that's going to transform legal education, and hopefully we'll get that integrated into our teaching and our scholarship as well.

[00:03:29] **Talley:** To help us think through some of these issues, we are thrilled to be joined by today's guest, Kareem Yusuf. Kareem is senior vice president of product management and growth for IBM Software, where he's worked for more than 25 years in a variety of senior leadership positions. Kareem has a Ph.D. from the University of Leeds focused on decision support systems for civil engineering construction, is an accomplished TED speaker, and is the author of a book, *Enterprise Messaging Using JMS and IBM WebSphere*. Kareem, welcome to the show.

[00:04:00] Yusuf: Hey, Eric. Thanks for having me. Hey, Dorothy.

[00:04:03] **Lund:** Hey. Welcome. Well, so let's start with the end of 2022. OpenAl publicly released its generative AI technology ChatGPT, leading to global hype around generative AI and its potential impacts. Now we're seeing tech companies rushing to develop and advance generative AI. In addition, governments are racing to adopt and regulate it, and business leaders and workers are all navigating how best to use it. Let's start with some basics here. What exactly is generative AI? What makes it different from other forms of AI? And what are some of the most exciting use cases for this technology in the corporate and legal sectors?

[00:04:43] **Yusuf:** Really good question. And some of us liken, if you want to be historical here, the advent of the OpenAl announcements, like the advent of the Netscape browser, right? And what that did for the internet, suddenly bring into the forefront of everybody's mind: What should we be doing with this technology, and what could it unlock? And I think that's kind of the place we found ourselves at the start of 2023 as a tech industry. And you're right, a lot has happened in 18 months. So, let's baseline a few things. When we think about AI, it has always been a journey of trying to train computers to detect and understand, if you liken the patterns in data to kind of

guides particular activities. And this has been a journey that's been going on for a very, very long time. If I dial back to the most recent times, and so I'm going to take you kind of to the '90s, the terms were expert systems, and a lot of the systems were rules based. And that meant you could say, you know, if you see this, do that, if you see this, do that. Obviously, you can immediately imagine the challenges in scaling that approach, but we built a lot of systems, indeed. When it really caught everybody's mind is what I would call the advent of machine learning. And this was the opportunity to really begin to take datasets, label them, and use those to drive an understanding of the patterns, leveraging essentially what people called neural nets. You would, to pick a classic internet example, take tons of pictures of cats, tons of pictures of dogs, label them as cats and dogs, and begin to send those images through the neural nets, adjusting weightings for the right matches. Extremely labor intensive, but even more importantly, required a lot of labeling of data, right? And so that became kind of the norm. And you would find many an application of machine learning still very active today. In fact, the thing that I think is important for everyone to understand is, we haven't abandoned prior techniques. Machine learning is still actually extremely good for statistical-type analysis, doing things like predictive maintenance or predictive analytics against any kind of dataset. So, for example, predicting when a machine might fail or what might happen next. What then happened was the state of technology moved forward to this notion of, we could achieve better results-even better results-with less labeling of data if we adopted this transformer technique that led to large language models. The downside was you needed even more data. And there's a reason why this was applied first in the realm of language: because what's the most vast publicly available data? The written word. And so that's how we actually began. And that kind of brings you up to date to kind of the technology evolution as it stands.

[00:07:48] **Lund:** Kareem, real quick, can I just jump in and ask, when you say the transformer model, can you explain what that exactly entails?

[00:07:56] Yusuf: It all comes down to, in my mind, math and weights. People just began to look at what are different ways you can help to derive and under patterns out of data. New techniques evolved. It required less labeling of the data, but it also required a vaster amount of data than even what we were doing in machine learning. And that's what led down to this notion of large language models. Large language models, which is what essentially drives the likes of what you referred to ChatGPT, really upped the game in two key ways. It brought to the fore this notion of conversational engagement with the computer-the chat-based interface. Because remember, it means the computer has to understand what it's being asked, then find the relevant information, and return it. And so, it really, I think, captured all our attentions because it really said you could pretty much elevate how easy it was for people to engage with this, with AI. And what you're seeing in terms of critical use cases, I'll put them in three big buckets. In the consumer realm, there's a lot of what I would call knowledge acquisition, right? Variations on search. Tell me about this. Help me understand this. What is the latest state of this or that? We've seen that move not just from the written word to now being extended a lot more into images: Give me an image that represents this or an image that shows that. And so, you've seen that play into the

creative realms with all the challenges I suspect we'll get to later related to that. In the business world, I will say it's really all centered around this notion of productivity. So, one of the most common scenarios is basically creating more capable customer service. Whether you read that as supporting customers of your product or internal-supporting employees or students in accessing services. We see a lot of work around HR, right? So, you have what you would consider a chatbot, where you're building a lot of capability around that with, you know, these LLMs to drive that. We have seen a lot of work flip into the areas of what I would call business process automation, or we like to refer to it as digital labor. Where it's not just about the customer service, but you're now actually trying to get the assistant, as it were, the AI assistant, to support in doing actual business process work. Staying with my HR example, if you think about basic customer service, please give me an employee verification letter. When you move into the realms of digital labor, support me in a job requisition process with all that entails in terms of communicating with multiple systems. The other last area is in the realm of application development itself. So, developers are seeing these large language models as key ways to accelerate the art and the work of building applications in terms of code generation, code explainability, understanding old apps that you want to modernize, or creating new. And I would say we still see everything kind of circulating around those three core areas.

[00:11:05] **Talley:** Let me push a little bit more on that, Kareem, because I will say that for years, you know, I've had kind of a side gig trying to do scholarship that was basically you could probably classify it as machine learning and law or machine learning applications of law. And for a lot of that time I spent pretty much pulling my hair out wishing I were a better Python programmer than I am. And one of the things that's pretty amazing about some of these recent applications is that I can sort of say, hey, write me some code that does the following thing. And it may not be perfect, but it's not too far off. I guess that boots up a much more, I don't know, in some ways, ominous, for some people, question about where are some of the use cases and applications of these new generative AI models that are extremely powerful, particularly amongst fields that have thus far been largely sort of resilient or resistant to automation waves, and, you know, the industrial revolution, like service industries, lawyers, accountants, radiologists. These are occupations that you go not only to college for, but to graduate schools. How do you sort of see it shaking down in terms of really powerful use cases?

[00:12:15] **Yusuf:** So I would kind of I actually think you have to look at the problem space differently. Creating applications like many of these activities is a science and an art. Yes, there is the explicit science of how you write the application. A lot of that becomes all of a sudden boilerplate and you know that, you know, this becomes standardized. Then there's how you want to apply yourself with that ingenious algorithm that does something distinct. You have to think about an algorithm. You write it in code, but you actually have to think about it and come up with what you want the algorithm to be. And so now you have tools that help to alleviate some of the more boilerplating and allow you to have more thinking time to think about crafting really interesting algorithms. And I think it just builds and builds and builds on.

[00:13:00] **Lund:** The way that you're framing this is it's just, you know, it's going to supercharge productivity, it's just going to be additive. But I've already spoken to CEOs of companies and directors who are seeing workforce reductions due to the use of AI. You know, we don't need a social media coordinator anymore because ChatGPT gives us our social media strategy. PricewaterhouseCoopers has expected that adoption of generative AI is going to lead to workforce reductions of about 5% this year. So, where is that coming from, and where should we be afraid?

[00:13:31] Yusuf: I think you should think about it as work shifts, right? Let's take some real-world examples. Call center for HR. You can have 700 people manning a line, or with the right kind of automation, drop that to 70 people. The point is, what happens, to as you raise it, the remaining-I can't even do the math in my head-630? Typically, what we're actually seeing, and even this is within our own company, you're able to shift to higher-value work. This has always been the promise and indeed the challenge of automation. It typically involves a level of reskilling. It does involve work shifting—people shift into higher-level work or other work if you're staying in the domain. But this has been the story of computing from the get-go. This is my point. I don't—I'm not by any means minimizing the fact that industrialization, automation changes the nature of work, but it has been changing the nature of work. And generally speaking, if you take a look at it and maybe consider me a techno optimist here, the way in which IT to information technology in particular, has helped and transformed the kind of work we do and our ability to therefore do more-interesting and high-value work. I don't think can be questioned. None of us will want to go back to the ways in which we were working 30 years ago. We like those productivity boosts, and I think generative AI is enabling that yet again.

[00:15:06] **Lund:** There's also another issue that I'd love to get your thoughts on. And this is about how creative professions may be affected. And right now, there's a lot of questions about the use of intellectual property training AI systems. And we are seeing this debate playing out in the lawsuit that *The New York Times* brought against OpenAI and Microsoft for alleged copyright infringement. And I know you and IBM have indemnified companies against intellectual property claims related to the generative AI models you're using, and you're actually even publishing the underlying data used to build your AI system. So, I'd love to hear your take on these concerns about the use of copyrighted materials by writers and other artists that are being used to train generative AI models. What do you think the solution is here? How important is transparency around training methodologies?

[00:15:59] **Yusuf:** I think it speaks to the ethos of the company. I work in IBM itself. Indeed, my personal ethos, as I think about this space, which, first of all, anchors on transparency. When you think about this notion of models and all this debate around copyright, right? And you guys are the lawyers, you will understand this a lot more than I do. It really anchors on, in my mind, at its core, fair use. So, the question is what data, and do you have the rights to those data? Our approach has been to show our models are built on datasets that we, you know, are readily available publicly or datasets, you know, we own and therefore have the right to use. This question of do you have the

right or are you giving appropriate recognition to those who have the right is a fundamental question of what I would call people's approach to building technology and doing business. It's not really the AI problem. It is the problem of you want to train a model, do you truly have your rights? Does it matter to you to think that that's an important thing to be able to do? And as you then follow from that-so that's the copyright side of things-you flow into the notion of the output. And so, we get the classic calls of wait, are you saying this model can literally reproduce the entire book that, you know, somebody wrote? That doesn't feel right, right? So those people then, you know, jump in. Or is the model giving output that seems offensive or unclear or you've heard the phrase "hallucinations," right. And so, people say, well, what are you protecting on the output? Because if I'm using this to support my customers, I don't want it to start saying some unexpected stuff to them. And so those become the two areas of where indemnification begins to play in: Do you understand what you used to train your model on, and can you stand by it? And how much support are you willing to give, under which conditions and which use cases around the output? That is where the conversations lie, and I think it still remains to be seen, this question on copyright, how far people can really go, right? That's why these suits, you know, being worked through will get clarity, I suspect, in the fullness of time. But I do think ethics and one kind of values is actually the word I'm looking for are important. And for us, it's important to be clear about what we use and how we use it, because we serve business. And businesses need to know what they're what they're using and what they're banking their business on. The stakes are high.

[00:18:35] **Talley:** Let me follow up on a point you were talking about with respect to hallucinations. There is a fairly large set of concerns about how the growing adoption of AI is going to pose really brand-new questions on how to prevent the production of fake content from things like deepfake audio and video to spam robocalls. Earlier this year, Meta called for companies across the tech industry to adopt technological standards for identifying and labeling artificial content that could be dangerous. What's your take on how tech companies can ensure that their AI systems are enhancing quality, not taking a turn toward the nefarious and sort of bolstering what we can depend on from the standpoint of online content?

[00:19:19] **Yusuf:** We should be looking at making it clear when is Al being used. You should know you're interacting with a, to use the phrase earlier, chatbot or Al assistant. You should know that this content is Al generated: robocalls or this. We've been talking to, you know, IVRs, you know, integrated voice response units for decades, as you call your bank, you know, to check stuff, but, you know you're speaking to a non-human. I think that's where it becomes really important. Once again, I'll come back to that transparency, because transparency builds trust and raises the right level of awareness that I think, as an overall thought guides—or should guide—how people operate. When you look at where these regulations are going, they all begin to center around what is the use case, how do you rate the use case, and what kind of checks and balances should be applied? And when you decompose the checks and balances, it really comes down to, are you consciously making decisions every step of the way that you, the the individual, you know the team as you deploy this technology, and do you have the

appropriate audit trail to show your considerations, right, on your checklist and your preparedness then to manage this technology kind of in the wild.

[00:20:35] **Talley:** Just last year, the European Union and 28 countries, which included the U.S., the U.K.—your home country—and even China signed the Bletchley Declaration, which seeks to increase international cooperation on AI safety. And the White House has doubled down with a rare collection of collaborations, even including the Chinese government, to work on safety of AI systems. Do you see AI-centered discussions actually becoming the diplomacy of the future? Are there ways to implement these sorts of agreements on an international scale? Or, on the other hand, would you anticipate that AI discussions may become a growing sense of tension between countries that just don't have the same vision about how to regulate the technology?

[00:21:20] **Yusuf:** If you take Japan as an example, as they've looked at the models they've been building, they've used very, you know, kind of open, I would say, perspectives around copyright in that regards right within the country and say, look, we're kind of elevating the technology and they're driving to what I would call very, very open usage, right? And others maybe not so much. I do think that ultimately this becomes one of those things where responsible business operating responsibly will kind of set some of the de facto norms for how these technologies play out. And you will see regions, obviously, that are more advanced from a government and policy perspective. Europe has been, by the way, I would have said we saw the same thing on sustainability, right? And they kind of set the tone. The EU set the tone for what a lot began to follow.

[00:22:13] **Lund:** Another topic here that you've spoken about is the intersection point between AI and goals towards net zero emissions. And so I'd love to hear you say a bit more about the role of AI in the net zero transition and the types of collaboration we might need to see to make this happen to support these sustainability initiatives. And I'll also note that there's sort of a double-edged sword here, which is that, you know, there's some research suggesting that AI is a very power-dependent industry that requires a lot of fossil fuel production in order to meet the needs. And so how does this all balance out for you?

[00:22:49] **Yusuf:** Al's superpower is decision support. If I think about it that way, it really is supposed to be enabling us to better make, be able to make decisions based upon more data than we could actually kind of get a handle on just ourselves. And I think in sustainability in particular, we're seeing this begin to play out around realms of what I would call climate prediction forecasting, things like biomass calculations. We did some work with NASA on a foundation model for geospatial data. And just because I haven't mentioned that term before, let me kind of give a quick connection point. The actual technology we're talking about are generally referred to as foundation models. These are the things that enable generative AI. A type of foundation model is a large language model which is focused on language. Models, the foundation models don't have to be focused on just language. So, we've built foundation models, for example,

for geospatial data, which is actually a lot more numerically centric and image centric. You can build a foundation model, and we have, around time series data. Think of that like, you know, data streaming off a machine as it works and what might you do with that and understanding that data. So we've done work on geospatial models with NASA that really have been about allowing people to more quickly gain insights from the vast dataset they have as you try to do region analytics. For example, determine the propensity for an area to have fires. And we've actually been doing some work under a heading of environmental intelligence where we've been leveraging those models to try and build simple access systems so that they can be integrated into decision support systems around climate, vegetation management. So, I do see, you know, this technology aiding us in sustainability efforts. Now, there is no question that foundation models, large language models in particular, they are compute intensive. And so, you do get into the question of, well, this is very energy intensive. And so, you come back to, well, how is that energy being produced? A lot of folks though, and companies in general, as we look at our energy consumption, continue to drive goals towards more renewable energy. Many big companies have renewable energy goals. I think you will find, if you take a close look at many of the big players as they're looking at their data center rollouts, linking more renewable energy resources into their data center consumption. And so, I think this just reinforces the need. The other side of the tech is, with also a lot of focus on being able to do more with less intensive models. Models are measured in terms of the number of tokens. So, you have a 3 billion [token] model, 30 billion token model, 2 trillion token model. And just assume the more tokens, the more intensive it is and expensive it is for that model to run. But if you can come up with smaller models focused to a particular task or domain-which is what we kind of fixate on back in IBM, because we think of those targeting particular business processes—you also can improve the computational footprint.

[00:26:10] **Talley:** What are your predictions as you look forward for the next wave of artificial intelligence? You know, there are at least some signs now that the explosive growth in foundation models in generative AI has started to level off a little bit, at least in large language models. That may be in part due to the fact that some users are starting to migrate to small language models, saying we have proprietary data that we want to carve ourselves off from these other models. How are the risks and the unknowns likely to grow as well?

[00:26:40] **Yusuf:** What I would expect to see, I think the democratization of the technology in its own right. Building large language models, foundation models from scratch is very intensive and very expensive, so only very few would do it, right, as a baseline. But do not underestimate the power of the open-source community. Some key practitioners of the art—think Meta, think IBM, even think Mistral—have all provided open-source models—Google has as well—which has created a whole industry around people beginning to craft models against tasks. Except there has been one very specific problem: There is, to date or very recently, not been an easy way for people to kind of craft these models and bring the knowledge back together. We introduced a technology as an open-source project just back in April, called Instruct Lab, that is really designed to enable that in the open community, where you can take a base model and we can all

contribute what we call skills and data to craft a model to a particular use case. I think that is important for two reasons. One, it actually begins to create this lingua franca, this nomenclature around, oh, I can improve and optimize a model guite easily and in a normal way to meet a particular use case: I want-skills and data added on to that base model, and I can do it in a way that is collaborative, so I can bring together people of like-minded people to bring it. And so, we put that in the open-source community. The power of that for customers, right, who are then looking to say, I want to train and modify models in my own right with my own data, they can do that in a way that's a lot more cleaner and capable of them saying, here is the data I added. This is the base model I was using, than ever before. And so, I think that becomes very important to making all of us true creators, as in true participants in this journey, versus it being controlled by a few who we can only access via some API on the web. Democratization means choice. So, we put it in the open world because that means everybody can contribute, everybody can evolve the technique, and there will be certain models people will want to build for certain situations that they also want to make open. If there's anything I want anyone to kind of take away if they listen to all of this: Just don't be afraid of the technology. The basic questions, the basic principles of living, running a business, operating-there is no difference. We've just applied yet another variable that we need to wrestle into and get a sense of. This, I actually think it's kind of exciting, but I accept the risk's a little scary. And I think we've kind of unpacked a lot of those. And I think transparency and appropriate governance and regulation to protect individuals and us taking personal responsibility to educate ourselves and understand what it is that we're dealing with at any point in time becomes the way in which you mitigate that.

[00:29:59] **Lund:** Kudos to you for really being able to walk us through this and help us understand all of these many different aspects of this technology and the legal issues and business issues that it raises for all of us.

[00:30:13] **Talley:** Yeah. Hear, hear. Kareem, thanks so much for joining us today. And I would just like to say for the record that none of the three of us sounded at all like Scarlett Johansson on this. [Laughter.] Our guest today was Kareem Yusuf of IBM. Kareem, thanks so much for being here.

[00:30:30] **Yusuf:** Thank you for having me.

[00:30:35] **Lund:** Join us next time for another episode of *Beyond Unprecedented*, and make sure to follow us on Apple, Spotify, or wherever you get your podcasts. Thanks so much for listening.

[00:30:46] **Talley:** *Beyond Unprecedented* is brought to you by Columbia Law School and the Ira M. Millstein Center for Global Markets and Corporate Ownership. This podcast is produced by the Office of Communications, Marketing, and Public Affairs at Columbia Law School. Our executive producer is Michael Patullo. Julie Godsoe, Cary Midland, and Martha Moore are producers. Editing and engineering by Jake Rosati. Production coordination by Elise Dunaway. Special thanks to Erica Mitnick Klein and Molly Calkins at the Millstein Center, with research assistance from Narda Vasquez Arancibia. If you like what you hear, please leave us a review on your podcast platform. If you're interested in learning more about law, the economy and society, visit us at <u>law.columbia.edu</u>, or follow us on Facebook, Twitter, and Instagram.