

Audio file

[Session 2 Eric Traugott moderator.m4a](#)

Transcript

Eric Traugott

Thanks, everyone. We're going to kick off, I think, the most exciting panel of the day, governance structures. We'll talk about what's working and what's not. None of our panelists have ever made a mistake, so they can't talk about the what's not working side of things, but maybe some anecdotes along the way we could talk about, see what experience they have. So Eric Traugott, I'm a managing director with Park Lexington. We've had the pleasure of working with Kudu doing some operations compliance and technology due diligence for some of their portfolio companies and other investments they were looking at. So we're pretty excited about this panel, obviously, as AI crosses over really all three of those different areas and provides a really nice perspective. We've worked on probably a dozen projects with them, and it's been really interesting to see kind of how Investment managers and wealth managers compare to the industry, but also compared to each other. And so, obviously we've seen a different evolution of systems and operational controls along the way. And the adoption of that is different at different firms. And I think we're going to see the same thing with AI. So kind of very excited for the next phase of investments in the infrastructure at asset managers and wealth managers. But Today we're here to talk about governance. And so maybe before we get into some of the questions, just thought I'd give the panelists a chance to introduce themselves and their firms, and maybe if you can talk a little bit about where you are in your AI implementation journey along the way.

Jeff Saginor

Thanks, Eric. My name is Jeff Saginor. I'm the Chief Technology Officer for EJP Capital. We're a multi-strat manager based outside of D.C. We are, I would say, intermediate amateur AI users. So we have a few tools that we've developed and kind of operationalized that are really enterprise-grade tools. And then we sort of have a skunkworks that is constantly assessing and determining if new tools are valuable and applicable to the business. So we've been doing that for about three years now.

Sumit Sahni

My name is Sumit. I'm a founding partner at Vaxa Partners. In a prior life, I was a strategy consultant. We founded an M&A boutique and strategic advisory firm that focuses on private markets players. And so the AI automation of private equity, private credit players is a massive topic. and as is all of the OPEX and the supervisory burden necessary, so as not to make mistakes and to kind of align and adhere with the kind of regulatory environment, which is going to get onerous, I think, as we go forward.

Patrick Hultgren

Hi, everyone. My name is Patrick Hultgren. I'm the Chief Information Security Officer at BTIG. So we're a trading firm. headquartered in New York, but with global offices, about 700 employees, about 70 people in our technology department. I like giving these stats up front so everybody knows exactly what the parameters are. For follow-on questions, we're in like year three of our AI journey. I double-checked an e-mail before I came up here, and it was from June, and it was from my legal team, and it said, confidential, Gen. AI. And that's when this all started, which will roll right into the governance part of it. We have probably similar to Jeff or like early phase two of what Sanjeev was talking about. We have a minimal viable product we're rolling out internally for our firm right now that is a combination of a bunch of tools that we developed in-house after doing the loopy loops around looking at kind of all the big box solutions. So that's where we are right now.

Eric Traugott

Thanks. So maybe we'll start simple. What's the bare minimum that you think firms need to have in place from a governance standpoint? And maybe it would also be helpful to understand if you had those things in place when you started or kind of where you developed the bare minimum along the way.

Jeff Saginor

Do you want to start, Jeff? Yeah, I'll start with this. And I love that level setting. We're about 100 folks and we have a very, very small IT team, a GC and a head of compliance. So we're working very lean. And I would say that we started thinking about this really because Sanjeev mentioned that there have been a couple technological revolutions over the last 20 years that are kind of analogous to AI. But I would say that really, This is a consumer-driven sort of revolution. Everyone in your business is using ChatGPT at home. They're hearing about it, especially the young people. And so it really requires, you know, the cloud transformation that I did 10 years ago, you convince your leadership, there's, you know, efficiencies, there's security, there's resiliencies, there's all these things no one notices in, you know, in the business necessarily. And these are things that your analysts, that your

ops people, that your legal team, They're coming every day to you and saying, I'm using these tools at home. They're letting me do these things better. I've got it in my pocket. I'm asking it questions. I'm using it like Google. Why aren't we doing something? So in my mind, that really required doing something different as opposed to just setting up the. a few policies that you work with legal and compliance on and maybe setting up a governance council and then moving very slowly, we really had to be more flexible and adaptable. So what did that mean for us specifically? It meant we set up an AI policy that was fairly broad. It involved not entering PII, it involved not entering MNPI, it involved not using AI to generate trade ideas and sort of move portfolio products. And then we said, fill out this form, submit it to us, and the team that is responsible for sort of driving the business, and that of course involves legal and technology and compliance as well, will determine if this is something that is appropriate for us. And then we will coalesce around those tools as we sort of find them and as they pop up. So we handled it that way. It required a couple things. It required deep buy-in from leadership because there are going to be many more questions that are not answerable by your risk team, which is kind of what I see the legal and the compliance side, and not answerable by your technology team. Yes, we could do it, but is the risk and the business aligned there? So we did that. And then we can talk about many, many kind of conclusions that we drew from that, but we found a couple tools that are really, really driving the business out of that process.

Sumit Sahni

So yeah, so kind of we come at it slightly differently. One of my observations, I think, in the industry is that AI governance has become kind of all referee, but no one actually playing the game. And so there are a whole load of folks in the kind of central functions, risk, legal, who are very concerned and very kind of proactive in trying to put controls in. But often I find there's a big vacuum in terms of defining what AI strategy is and how it's going to redefine and transform not only the business but also the sector. And so I think there needs to be a creative tension between the art of the possible and the kind of risk and compliance view. And I think as an organization, if you, a lot of GPs, particularly in private equity, are quite proactive in making their portfolio companies develop transformational value creation plans, but very often don't have their own. And so it's really important, I think, from a governance perspective, the GP to have a VCP where AI is singled out as a transformation lever, where the AI transformation is detailed and valued both in terms of earnings and sort of multiple impact. And then all of the control mechanisms are then built in to the execution of the VCP. And so then you get into PII, legal risks. There are also a whole load of financial risks as well. And I don't know if you guys are familiar with the EU AI Act and the kind of Sr. 11-7, all of the OCC Fed oversight on model risk governance. All of those things have to be taken into account. Because I think it's inevitable that if you're an investment

manager, 80, 90% of what you do is going to get automated and based on agentic and kind of Gen. AI. There's a whole can of worms then as to how you organize the governance structures to make sure that not only do you play the game and win, but also that you have that kind of refereeing capability. And I see a whole spectrum of maturity in terms of driving value creation plan discipline and transformation, but then also making sure that you adhere of regulation and don't make kind of rookie errors in implementation. And we can kind of talk further about that.

Patrick Hultgren

Yeah, so we're, like I said previously, we're a trading firm. And downtime means you get fired. So we're very, very conservative with anything we're putting into production into our environment for all of our legacy systems that have been built forever and have been hugely successful and are totally bespoke and arcane all at the same time. So our journey through AI started, which is relevant for governance with, like I said, our legal team coming to everybody and saying, We got this. We, the brilliant legal minds that have served you so well for the last 20 years are here to save you from AI, which people are going to use to be lazy, not work, and put like stupid haircuts on their dogs. Like that's what this is going to be, right? Okay, I'm not sure, but like that's great. We're glad you guys are proactive about it. And so, you know, that took a year for everybody to unwind themselves from. And eventually, We have really good executive leadership on this who have a buy-in and have a desire, just like Sanjeev was talking about, to move this forward and have tools that everybody at the company can use. So our biggest mistake, just from a governance structure perspective, was letting anybody who would be on the risk side or any specific shop. It could have been me, it could have been security, right? Saying, oh, I got this because I know security and I want to make sure the firm is secure, right? So letting anybody in that specificity run with it was not good, right? It does have to have, you need your executive leadership and then you need, your teamwork underneath that. And then I think that obviously, whoever your technologists are, your technology people have to be investigating these things and being proactive. We have a developer who, I mean, completely ran with this. We all know him. He's great. He's been at the company for a long time, and we just kind of gave him some left and right limits to be the driver of the product we're coming up with for the firm for everybody to use. And then once you have your technologists and your business leadership pushing in the right direction, I think the governance structures kind of took care of themselves a little bit. And then everybody gets in their happy place and sits in their lanes and they say, I'm security, I'm here to give you perspective, I'm legal, I'm here to give you your perspective, so on and so forth. So it's contextual within the firm that you're at, what your labor is like and who's in those management positions. But that's like my number one caveat, is that anybody who says,

hey, we got this, we're going to run with this thing, that's not going to work. It's got to be a total buy-in. Everybody jumps in together. You get your leadership giving you your left and right limits, and then everybody's working together to move forward. And as you build and as you grow, the products will present itself, the use cases will present themselves, and then you can augment your structures for governance going forward. How often do we need to meet? How many people really need to be in on this? and you can start delegating and all that fun stuff and be off and running.

Eric Traugott

So we talked about policies and procedures. We talked about a governance council and an overall vision. I mean, do you guys feel that the governance structure needs to come first, or is this more of a trial and error with the AI tools themselves, and then governance can come after you figure a little bit of it out?

Patrick Hultgren

It's going to be leadership dependent, but you need that. So like, you know, you need that connection. human, professional, otherwise, with your leadership to know what their end state is. I would think you got to figure that out a little bit. And I think, Jeff, you were probably speaking to that earlier when you were saying, this is a little different. Like it's going to be contextual because we're dealing with a whole new space here that we haven't seen before.

Jeff Saginor

Yeah, I can add some kind of real world feedback there. And I think the reason that it's different is we as technologists, and certainly as the executive leadership of our firms, typically think about any changes as top down. We think of them as strategic, that we are in the position to make those decisions and understand the business. And I really, truly feel this. This is an opportunity to have some of these really, really powerful and transformational ideas come from somewhere else in the business. And whether that is your frontline marketing person that knows their job better than anybody else, or your fund accountant that knows where all those pain points are, or some analyst that is doing pivot tables in Excel and could never get that out of their head to a developer or to a business analyst to help them understand the process. you're democratizing that. And there has to be a flexible process that the leadership and that the technologists that are going to end up on the security side and on the development side and on the sort of Dr. and resiliency side have to manage this, there has to be a process to get those ideas to us. So what does that look like specifically for us? It is a Microsoft Forms form that we are extremely responsive at

integrating into our workflow. So it's a form that goes to AI at ejfcap.com, it goes to me, it goes to our head of compliance, and it goes to a couple other people, and we have a meeting every single week to discuss these use cases, and we action them very quickly, whether that's The Zoom AI summary that is very effective, but comes with compliance risks as far as electronic records retention and things like that, or 100 other things, and the point is that your users are already using tools that are starting to integrate AI you may not know about them. They may be CRMs that start pushing AI features. They may be valuation engines that start to push AI features. And you need a process to be able to understand how that's going to affect your business. And your frontline folks are the people that will understand that the best. So that is kind of how we, why we set that up the way we did.

Eric Traugott

I mean, I could.

Sumit Sahni

Yeah, I was just going to add to that. I think that's really important is to not, obviously we serve larger, more structured players. And often that kind of bottom-up democratization suffers. But as well as having that, and I'll give you, this is my bias, you've got to have a point of view as to in five and 10 years, how will the end-to-end value chain and industry be completely transformed by AI? Of that, what is tactical that you can do now? And then you also have to mark to market, because a lot of the best practices in the art of the possible are coming from maybe direct competitors, but often kind of adjacencies. And so this is why, so for me, the sort of the governance conundrum is to create a non-bureaucratic way of capturing all of that and start to execute whilst also then setting up some of the control mechanisms. and having the control mechanisms embedded in the development of that plan. Because I think if you don't, if you're not proactive, you will get left behind. And really the litmus test question is, you need to be able to tell your board how much OPEX and how much growth AI is going to deliver on a quarterly basis. Because if you don't have that, you know, as a GP and then also with portfolio companies, it's the money that talks, right? This is not a VC environment where we can do hand waving and hope for a big exit. You've got to have a transformation plan and all of the controls as part of your kind of governance regime.

Eric Traugott

I mean, I would imagine if you have this led by legal and compliance, they're going to tell you all the reasons why you can't do something. If you have it led by the business, then hopefully legal and compliance are there to support you and protect you.

Sumit Sahni

But often the business doesn't know, right? I think the keynote speaker was sort of alluding to this, is that the hats that you have to wear to be able to answer this question, how is AI going to transform the business? You need to be multi-hatted. And there isn't one single person who has all of that insight. And often the answers are not, even if you did have the hats, it's going to be BlackRock who are doing something really cool, right, in their indexing desk. Or it's going to be, you know, Carlisle. I mean, you know, I have a view on some of the best practices, but there are players out there who are doing little slivers of the right thing, both in terms of the strategy, but also in terms of the control. And you've got to learn fast and continuously mark yourself to mark it.

Patrick Hultgren

Yeah, I'll follow up there. One of the things that was hard to do as we were being crushed by our internal bureaucracy was just come up with wins, like good use cases for the firm. And we should have been more aggressive in technology of coming up with that. Like we were not good enough about identifying really good use cases that we could use. that we could then, look at all the variables and articulate to leadership and say, hey, this is real. This is something we already did internally. This is a win to give the people who are doing their natural job, you know, from a risk or a policy perspective in the abstract, getting something to really sing into and say, actually, no, it's not, you know, it's not people being lazy. it's XYZ and these things. And so once we could start to ground, once we could bring those wins to that kind of governance structure, then they can orient and we could move forward. But super, super important, like you have to battle for those. Like you're going to be, you know, the governance is going to be overbearing to begin with because there's not enough product or content for them to sink into. So you've got to come up with something and give them that, and then you can start to work and use that as a reference point for going forward.

Eric Traugott

with these councils, you have business, you have technology, you have legal compliance. Who makes the decision in the end?

Patrick Hultgren

Executive leadership. It's you bring them the decision point and you fight it all out and then you bring it to them and they're going to say yay or nay on where they want to go forward.

Jeff Saginor

Yeah, I agree. And as I said, I think it is not like the, it is like in so many ways traditional technology transformation sort of projects. But in this way it is not because typically that's executive sign off and then you implement and you're successful or not. And that is the metric and you do what you said you were gonna do and everyone's happy or not. This really does involve a level of executive interaction. They have, there are going to be so many of these tiebreakers that are just like, is this too risky to do? Is this a benefit? We're thinking about hiring another person. Do we go all in on this right now and not hire that person? Those are the kind of decisions that have to be made literally on a weekly basis sometimes for some of these tools. And so you really have to let your leadership know that that's what they're getting into. I think it's certainly worth it.

Sumit Sahni

Yeah, I mean, I think the thing that, you know, I've actually experienced and helped lead is a mashup of a typical transformation program, VCP, but you run it in agile. So it's really the application of sort of agile delivery mindset and disciplines, but to a strategic transformation program. And that really works to get this sort of balance between top-down, multi-quarter view, and then bottom-up, quick wins, democratization, and a multifunctional view of the trade-offs between the upside and some of the value at risk if you screw it up. If this is better for the compliance panel to tell me, I'm curious what your regulator's reaction to this and what are you hearing from them?

Patrick Hultgren

What are you hearing from them and what's their temperature generally? I had to fill out a question here yesterday. So I'm sharp on this. It was actually insurance related. And holy smokes, it was all the normal questions that I would see. And then there were 10 just very generic, no choices, just big boxes. How are you using AI? What is your AI governance structure? So it's early days as they're catching up, at least for me, in terms of what they're looking for. But now they're asking in a generality, which means that it'll be like in two years, they'll be very specific about it. So that's about how much lead time we probably have before the regulatory crackdown comes.

Sumit Sahni

Yeah, what I've seen is SEC, OCC, and then the Gestapo in the EU are basically, they're fishing. They send out these surveys, it then hits typically legal or risk, and then the chaps in legal and risk look at it and go, I don't know what to do with this. They then call counsel,

send it to the law firm, and then the law firms sort of help fill it out. And then I think from a regulatory perspective, I think it's absolutely critical and it's not, obviously the US is a much more laissez-faire regime. But when I look at the EU AI Act, I don't know, are you guys familiar with that? Does it come on your regime? It's essentially a direct analog of the SR117 approach to model governance, risk, and review. So this is making sure that when you deploy AI, particularly in terms of asset allocation, asset management, trading, et cetera, that there is transparency, that there isn't bias, that there's replicability, that there isn't black box, that it's stress test, that it's stress tested, it's explainable, it's documented, that there's version control, right? All of these things that you see in model governance within risk management, this is kind of, I think, where the regulators are going. And so when you deploy, and it's just, I mean, it's just hygiene, right, to think of it in those terms. So wherever you apply AI, just make sure that you're kind of familiar with that SR11-7 regime and that you build in the controls as part of the development process so that you don't come a cropper in two years when it gets hardened into legislation. Now in Europe, for high-risk AI applications, I think as of last month, it's already a legal requirement. So if you have a business that does anything in Europe, any client, or serves any investor or client in Europe, you have to be adherent already. So that's a bad side. That's the sort of the control side of things. And so this is a problem. A lot of the people who think about that, and PII and legal, et cetera, they're actually running the show, which is why I'm really emphasizing this. You need a value creation plan. You need your C-suite and your investors and your board to be really excited about the value creation opportunities. And you need to be really specific so that you can then, and then you think about these kind of controls issues.

Jeff Saginor

I will just say that we have not had the SEC or the NFA ask about AI specifically, but LPs ask all the time in investor due diligence. And right now, to Patrick's point, they're asking, do you have an AI policy? Are you using AI? How are you using it? Maybe casually in a call. But next year, it will be, tell me the types of data that you're generating from AI, tell me the types of inputs and outputs, and then the next year, tell me the roles that you're, you know, with agentic AI. So that, I absolutely see that coming down the pike, and we've seen it just in the last six months.

Eric Traugott

Yeah, I mean, my feeling is they're asking those open-ended questions, because in the end, they don't really know what to ask for. That's right, yeah, they'll use your answers to come up with the questions for next year, but.

Patrick Hultgren

We don't store any data, right?

Eric Traugott

But I guess going back to something like the council, when you think of who sits on that and the expertise, do you feel that the... folks that are on that council need to have AI expertise in order to do that? Because a lot of people just don't have it. They talk about it, but they don't really have the in-depth knowledge. So how do you balance the lack of real in-depth industry knowledge in this space with having a council that should have knowledge on it?

Jeff Saginor

I will give sort of a real-world example, and I'm not sure this is going to answer the question. We, so we, this is sort of speaks to the flexibility and the agility of having to make these decisions. But we use ChatGPT Enterprise. We have an Enterprise license. We allow certain firm data to be sort of analyzed and generated from ChatGPT. And recently, ChatGPT enabled a SharePoint connector. Now, this could be a highly valuable tool for an organization like ours that we tried in 2022 to build out, OpenAI and Azure connector to our data lake and the level of kind of industry expertise in rags and understanding how, vector databases work. It was completely kind of something that we could not handle. And this solves that problem. but it also raises so many questions about the governance of this data. Had we anticipated that SharePoint would become essentially a data lake that everyone could access with a simple query in ChatGPT? What's the latest compensation figures, for instance? Are we confident that data was permissioned properly, protected in a way that it wouldn't go out the door? All of those things. That's where this council really comes into, I think, its value because your legal and your compliance team are asking that question, and then your technology team is answering that question, and none of this has to do with AI. The AI is the interface and the business driver, but you have the expertise on the council to get everybody comfortable, and then you decide, in our case, yeah, we feel very comfortable enabling this and searching side letters in two seconds instead of using a \$50,000 a year product to do it. So this is sort of a real world scenario, but something that I think this council, this type of council is very, very well suited for.

Eric Traugott

You guys think that AI experts need to be on these councils?

Sumit Sahni

Well, actually, I think this idea of the amorphous AI council, kind of a smoke-filled room with every discipline kind of hanging out, not really accountable to any value creation

metrics, is a kind of scary one. And so when you look at the big banks, again, the model governance is really going to be the precedent for how to manage AI. What I see some of the big guys doing is they have, first of all, like you have a risk register, you have an AI register. It's kind of a list of all of the AI use cases. So if you're going to want to do anything with AI, it's got to go on that register. And then the register is systematically assessed. It's an intake process. For, the metadata is created, what is it, what's the use case, et cetera, et cetera. And then what you want to do is give, have a, sorry, have not necessarily a council, but have a cross-functional group, a workflow where people actually analyze the register and assess it for certain risks. And then, you know, if it's PII or if it's operational risk or whatever, you can just send it to the right person so that the analysis is done. Now, That should trap kind of 60% of the unnecessary meetings and kind of powwows that would take place as part of a sort of amorphous AI council. The question then is, I think there are governance mechanisms and different institutions have taken different routes. So the less mature will have a head of AI and some kind of governance mechanism maybe sitting in legal. So I was just making a note of this. I think Carlisle have done that. They've got some bod in legal, they're the titular head of AI, and then Carlisle doesn't have a register as such, but they kind of run a kind of a council mechanism. There are others, I think CVC have AI is run, so on their board, in their risk committee, they have a subcommittee for AI governance. And so they have somebody kind of who sits there who runs AI governance, but it's very operationally risk focused. So it's quite narrow. And so there are various permutations. I've seen even an HR version because the issue is it's kind of, you know, it's basically managing the permissioning of internal data that's being put into an LLM. So it's about user access rights. who can see what kind of data. And then I think the big boy, the big players, the big banks, the GCIPs, I think have more complicated, but I think really good, and you guys could sort of, I think that's, if you're going to mark yourself to market, it's creating a more agile version, say, of what JP Morgan have done, where they've AI governance sits in a new organization they created called, they have a chief data and analytics officer who kind of sits at the center and is basically their responsibility is to create the strategic plan. And they have mini-me sitting in each of the lines of business. And then they kind of coordinate. So there's a whole, I mean, I can go on about this for hours. There are multiple permutations. But to cut through all of this, you need a VCP. You need a register of AI deployments, and then you need a couple of functional people to collaborate with so that you can start implementing things.

Patrick Hultgren

Eric, I have a quick question for Sumit. You were talking about quantifying the value of AI to an organization to essentially get senior management's attention.

Sumit Sahni

Can you talk about?

Patrick Hultgren

How you would do that or what metrics you use?

Sumit Sahni

Yeah, so, there are different horizons of impact. So there's a sort of tactical augmentation. So, meeting notes, automation, et cetera, et cetera, which is much more conducive to sort of democratic ideation, but needs to be managed and sequenced properly. I think essentially what you need to do is there's a little bit of strategic analysis about how do we use AI to automate our current processes, and then how do we actually transform them to be next Gen. So if you were alive in the 90s, do you remember business process re-engineering, right? It's basically that now, but with a much bigger bazooka, right? Nothing has changed. So it's a combination of process mapping, business model, blueprint definition, and then systematically looking at best practices, looking at the technology, and making bets on how your process environment can be transformed. And then there's a third horizon, which is basically industry transformation. So if you're in private credit, and if you're Apollo, you know, or Aries, private credit is going to get totally transformed. We're going to go from pools of risk to risk, tokenized risk at the individual. You'll be able to get real-time data on every obligor. It'll be connected to other real-time data. It could be collateralized. If it's collateralized, you could get real-time data, again, tokenized through the blockchain on their salary, if it's a business, on the cash flows, through maybe a credit card receipt. I mean, we've done looks into all of this. So there are some Star Wars versions of how AI will transform each of the asset classes, real estate, again. So you need a point of view about the tactical, the process and business transformation, and then the industry transformation. And then you define which initiatives you want to launch. And then it's just like a re-engineering program from back in the day. There's a current state, there's an end state, there are interim states, there are initiatives, you need CapEx, you need OpEx. And then there are P&L impacts and growth impacts. Then you prioritize, you sequence, you create a business case, you commit to the financial outcomes. And then in parallel, you do all of the risk assessment around, okay, for these initiatives, okay, what are the OCC or SEC going to say? And how do we preempt that? Oh, does this impact Europe at all? What are our adherence requirements for the EU AI Act, et cetera, et cetera, PII? So this is a big, massive opportunity, but with, a lot of work to get it right. And it may be that fintech, AI natives may just be more advantaged because they don't have to worry about all of this

transformation stuff, right? They can just do it natively, and if they fail, they die, and if they succeed, then they get big market caps. But that's, again, VC land. So, sorry, yeah.

Eric Traugott

I mean, oh, yes, Brian.

Audience question

Can I just ask, so we're a 20-person firm, investment manager, so we have a lot of analysts and PMs, and they're doing their own work. My question is, how do we, how do I make sure they're, if I secure the documents and make sure that they're not able to put PII in these systems, is that enough? and are the privacy settings and the tools now are reliable? Like, can I trust those, ChatGPT, turn off sharing and, copilot, whatever. So, how do I, make sure where I want people to experiment? I want people to use these different tools, but as long as I can secure the documents, great, but then what from there? What do we do?

Jeff Saginor

I will speak to this because I have to sort of wrestle it every day. So in my mind, these are consumer tools. Even, like Microsoft Copilot, which is, ostensibly a ChatGPT wrapper with purview behind it or whatever it's going to be. They do not by default, ChatGPT is a great example. They are not a governance first product. So even ChatGPT Enterprise, they're not training the model on your data, but they just last week introduced role-based access. That is insane. You know what I mean? You can't even, you can't divide up your entire organization based on what tools they can use within the platform. So you have to have, this kind of goes back to the fundamentals, but you have to have these fundamental sort of parts of the business, regardless of what these tools are. So in our case, that was a really strong DLP program. We wanted to be able to tell what was going out over ChatGPT, whether we were trying to control for it before it got to ChatGPT or not. And that's not easy to do. We use a tool called Cyber Haven, which is really, really good at understanding what's actually happening in the browser as opposed to at the data document level. And so it's doing things like that. So really understanding what the risk is. So hey, something could actually get out of these tools and where can I put safeguards in to stop those things?

Eric Traugott

Any other questions? I have one more, but I'll open it up if anyone else. All right. I mean, maybe my last question, you know, You mentioned this amorphous AI. I think people think of this as a really big deal, but I know, Jeff, you kind of mentioned earlier, like, this is a new technology, a very impactful new technology, but we've gone through cycles of new

technology before. So are we overblowing this in any way, or should we just have some similar governance and oversight as what we've done in previous technologies? Like, how is this one different? I'll go to Patrick first.

Patrick Hultgren

I'm going to give the classic it depends, which is the answer to every question. Depends on your business, depends on your organization, depends on like what this gentleman was talking about here, like the amount of labor that you could appropriate to something like this. But pretty early on, you know, you need to identify how that's going to be specific to your organization. So in our case, it's like, do you need to hire an AI engineer? Basically, it's like one of your first questions at organization my size. And that's like an investigation, which is a lot of talking to peer groups, you know, doing some research, doing all those things, talking to everybody who would be on that council, okay, from your lane, you know, where do you see this hitting? in trying to identify that. And for some bigger organizations, obviously, all the big banks answer is like, of course. So they jump into all that. We were kind of in the middle. So it's like we had some talented people in technology, and once we were able to get up and running and kind of get some runway, we were able to do it in-house. But you'll see that pretty early on, and then it's just, you know, what's the most cost-effective way to come up with solutions? And, you know, just been talking about you should really challenge, somebody internally to see if they can knock it out. Like there is a absolute, career ahead of somebody who can solve this internally at your firm. And they're going to have that every step of the way. They're going to be into the ground floor on the technology. They're going to understand all the process. They're going to want a pay raise, you know, they're going to, but you know, they're going to have that incentive to do all those things. So it's going to be organization dependent at the end of the day. And that just, you know, requires the faithfulness of the Governance Council to rock out to. And then we're done. We have one question here. Yeah, go ahead.

Audience Question

It's really for BTIG. So with trading and a lot of dealing, there used to be kind of old-fashioned dealer to dealer clients and then became electronic trading, which has matured over the past years. How does AI change or evolve from even that? Like where do you see your business going in terms of the implementation, in terms of your day-to-day trading or dealing or that part of the business?

Patrick Hultgren

That's a great question. I'm so not competent enough to answer it. So we are rolling out a, essentially a UI to use internally for everybody as a resource. Like Sanjeev, I thought his way of articulating it was really elegant. We're in that like that second phase where we're putting out a product that everybody can use that's all, you know, it's a dice wrapper, it's all locked up. We can see everything going on in it. What that's going to do to our trading platforms, I don't know. That's definitely going to be on those people in electronic trading. and how they build those systems going forward. So we're in the phase where we're rolling out that's something that everybody can use within our firm for AI, for internal things. How that gets put into the systems is going to be the next step. And as just security guy over here, I'm not quite sure. But I eagerly wait to see what those guys come up with.

Sumit Sahni

I think BlackRock are pretty hard core. And if you see what they're doing, it's pretty awesome. And what they've published is fairly dated now, so they've probably gone miles and streets ahead. I have some of the information if you want.

Eric Traugott

Just maybe in the last one minute real quick, just because it's what's not working as well. Is there anything that you guys implemented early on or that you've seen implemented that ended up having to be rolled back? Any policies and procedures that you maybe did incorrectly or as you learn more you realize weren't needed?

Jeff Saginor

I will just say that there are still many applications that AI is just not great at, and it will be. I am very bullish on sort of the transformational capacity of it. I think I'm less bullish on the two to... three-year transformational capacity for some things. I'll just give an example. Something that we do is a lot of loan document analysis, and these loans come from small banks, and they're not standardized at all. And we have always had a really, really hard time of getting this scan of a picture of a 20-year-old Excel spreadsheet that you can't read the numbers into any kind of data warehouse where we can actually do valuable analytics on it. And they come from all different banks. The loans could be two loans or 2 million. This was a real promising thing for me. AI is very good at OCR. It's very good at pulling out sort of an image and turning that into a table. And we can do that in many aspects. it doesn't do this well. It's just like a person looking at a three and seeing an 8. And if that's the case, we can't trust that. And that's what I'm sort of trying to drive home, this risk in the business versus the capability. I have no doubt it will get there, but it's not going to get there tomorrow for us, and therefore we can't change the business quite yet. But I want to see.

Sumit Sahni

Yeah, I mean, you know, we're a smaller firm. We were very bullish on using LLMs as a substitute for flex analyst capacity. And I don't trust the research that the LLMs do. They interpolate and hallucinate A lot. And so we have to be really, really careful in making sure we get the sources for anything that they claim to be true, because literally 60% is BS. And I've been using AI a lot in the past year. And so, I've spent 30 years as an analyst all the way up to a senior partner. And literally, you have to kick the LLM's *** because they just, the information it gives you is just wrong. So you have to be very, very wary when using AI tools for research exercises.

Patrick Hultgren

Yeah, in line with that, our first block of training that is coming out with our product is all around prompting. That's it. It's like we're trying to educate ourselves and everybody at the firm, this is how you write a prompt to get to where you want to. And like I said previously, we have a couple use cases that are really incredible, that people came up with on your own, and you just sit there and you go, holy smokes, that must have taken you forever. And they go, well, no, I just use the AI tool. What's that? Our prompts? Well, we took our... Oh, The good one is on our security team, me and my partner, he's been at this firm for 20-something years filling out DDQs. And he took all of his stock language in all of his DDQ information, which he would go manually cut and paste, put it into the models, built a prompt for it. So now I can answer all those, even though I haven't been at the firm for 25 years. And it's like, job done. And then, of course, like you're saying, we validate it. We don't just send it out. But that whole process can save literally days and weeks in terms of being able to come up with those answers quickly.

Eric Traugott

So I think we're out of time. We're maybe a little over. Sorry about that. But thank you to our panelists. Really appreciate it.