

# Audio file

[Session 6 Niall moderator.m4a](#)

## Transcript

### Niall Boland

All set. so I was reminded earlier that the only thing that doesn't change is change itself, because I was hearing a Reagan quote, which is, everything is changing too fast for me. The future keeps arriving, and nobody seems to be telling anybody about it, telling each other about it. So I'm glad that we're all here telling each other about it, because it's been fast. Nick and Katie and team, thank you for inviting me back again this year. I think it's pretty remarkable to see how much of what we talked about last year at a very high level of abstraction has actually kind of borne out. Nonetheless, it's still pretty hard to navigate. Nick used a phrase that kind of resonated with me, which was engaging with AI right now is a bit like you can go to the side of the pool. And one of the reasons that that's scary is it's the unknown is scary. And so it's useful to actually try and get a picture of what that pool looks like right now, because by God, it has changed fast, and in particular, where the deep end of the pool is. So with that in mind, I'm really glad to have a panel here of experts on the topic, which we think is the biggest driver of value for investment firms, which is the actual investment decision-making process itself, and the investment performance and scalability that goes with that. So with that, I'd like to introduce each of the team or each of the panel to introduce themselves. But I'd like them specifically to explain what the product is that they offer, what the use case is, but in particular to tie it to the business outcomes that they're trying to impact for their clients. And then where any of you have actual investment experience, investment decision-making experience, particularly interested to understand how your experience in implementing AI has changed how you see that role going forward. So we'll start with you, Bernie. Thank you for making it.

### Bernie George

Yeah, I'm Bernard George, co-founder and CEO at Tornado. Tornado is an AI synthesis layer for sophisticated financial institutions, asset managers, and hedge funds. A bit about my background, I started my career at JP Morgan. I was a proprietary trader there. trading corporate credit, cap structure, our equities, options, investing the firm's capital, internal hedge fund. Also was at Credit Suisse as a prop trader. And most recently before starting Tornado, I was a PM and on the investment committee at a \$13 billion fund that was spun

out of Barclays and then also built Managed portfolios there, so managed fundamental and quant portfolios there, and also built firm-wide portfolio optimization and Bayesian signal aggregation platforms. So as Niall Stewart pointed out, I think a lot of the opportunity in AI really comes at the intersection of deep domain expertise on how alpha is created and deep technological expertise on what is possible and how high that ceiling is. So what we've built with Tornado is essentially an orchestration layer which ingests KSQs, earnings calls, analyst days alongside your own internal data rooms, so private credit data rooms, decks, models, transcripts, loan docs, all of that. And then on top of that allows you to orchestrate essentially human-like analyst associate VP quality workflows. So And the ROI is amazing. It's essentially human-level work, faster, cheaper, better.

## Niall Boland

Thank you, Bernie. Over to you, Don.

## Don Muir

Sure thing. Hi, everyone. Great to be here. By way of introduction, my name is Don Muir. I'm the CEO and founder of F2, the AI platform for private markets investors. Before founding F2, I cut my teeth in traditional investing. I worked at two global private equity funds, Onyx Partners, which is traded on TSX on their \$7 billion flagship fund, and Apollo Global on the buyout team, investing in public to privates and sponsor to sponsor transactions. I went to Stanford Business School, which is where I founded my first company called Arc, which is a debt capital marketplace in the lower middle market. So on the supply side, we have 300 private credit funds and commercial lenders. On the demand side, we have thousands of lower middle market companies. And then we make the market between the two. When ChatGPT first dropped in 2022, I saw an opportunity to move up market by automating the debt placement experience within my B2B debt marketplace. And so I actually incubated what is now F2 as the technology to operationalize the debt placement experience, screening, underwriting due diligence of private credit transactions for the 300 funds and banks on the supply side of the Arc Capital Markets platform. Yesterday, we announced the official spin-out and capitalization of F2 as a standalone business alongside 10 million of equity from 50 of Arc's inside investors. Now we're serving about 50 private credit funds and private equity funds, automating core screening, underwriting, and portfolio monitoring workflows with our proprietary agentic system. So looking forward to meeting everyone and chatting about AI today.

## Benji Cohen

Congratulations, Don. I'm Benji Cohen. I am co-founder and CEO of Camber Credit, which is a vertically integrated private credit fund focused exclusively on the \$1.5 trillion U.S. auto market. Before this, I started my career in banking, as many people do in this industry, and then moved on to FinTech. And so I am actually a technologist in fund manager's clothing, which I think is what brought me here. T-Rex was a, or still is, a financial technology company that was focused on the structured finance market and found our way into private credit. We provided cash flow modeling, data management, and managed data service, and ultimately reporting. And this modularity made us effectively an operating system for private credit. So the Blackstones of the world pulled T-Rex in when they said, look, we've got 300 billion dollars of assets under management that we need to deploy. But the old operating systems that we had for credit, namely CLOs and public markets, don't work for this. And so we need to build within this Blackstone civilization something that can monitor all manner of these investments and allow us to deploy that. And we ultimately got to the place where we were fully automating \$55 billion of AUM every single day for Blackstone and ultimately for their LPs. I sold the company a year ago and was approached by one of our customers who said, look, this market, the auto loan market that I described, I've operated in before, but we need an AI native approach. And so he said, you know, you haven't been a fund manager before, but you have been a technologist, and that is what is going to allow us to operate in this market at scale. especially it being the last large fragmented credit market where others have not been able to do it. And so we are now building this and I would say have built successfully so far this AI native fund manager for this particular market.

## Niall Boland

Cameron.

## Cameron Ladd

I'm Cameron Ladd. I am CTO and principal of Sage Advisory Services, we're a roughly \$30 billion fixed income money manager. Kind of our AI footprint, we've been focusing and consolidating our data, really focusing on the data itself, and then now building out kind of structured AI, limited use, limiting basically which models that are I guess our employees can use. Right now we're shut down to just basically our Azure models and OpenAI models. We leverage Microsoft Azure, ChatGPT, and Copilot. And so we pretty much blocked all the Geminis. That's our current kind of AI footprint. We're trying, the end business goal for portfolio management is to really have our portfolio managers interact with the data, like canonical gold copy data, and be able to visualize on the fly different ways in slicing and

licing the data, really focusing on kind of a business intelligence and portfolio intelligence layer. And so that's been kind of our main focus as far as AI and portfolio management.

## Niall Boland

Great, thank you very much, Cameron. And for what it's worth, my background is I'm a co-founder of Clear Macro, and we work with large asset managers and hedge funds to drive their investment process and improve performance and improve the non-linear function of their business models effectively. I have one or two observations that we've had in the last year since I last spoke to you that I'd like to use by way of framing, and then pose the panel a few questions. So the few observations I have are that a lot has changed in the last year. What we find is often misunderstood is that there are a few different curves that one needs to think about differently. There's the curve of the foundational models, which has been extremely steep, and it's fairly well understood. There's also the curve of the agentic tooling layer that one can put on top of the LLMs. We think that's very poorly understood and is very powerful if done well. And then there's the adoption curves. And most of the surveys we've seen are actually showing up to 90% failure rate in early pilots in industrial settings. And we think part of that is a lack of clarity on vision and some of it is a lack of clarity on the agentic tooling layer. So that's observation #1. Observation #2 is in the investment process, we often see a blurring of distinction between what is the level of leverage being applied in the investment process, whether it's an aggregation tool all the way out to a critical reasoning tool where it can become an autonomous virtual research and investment decision-making team. So there's a whole continuum, and we usually find that most people aren't quite clear with themselves what it is they're trying to achieve. And then the third element, which is kind of part of the debate here, is the promise for AI in the investment research proposition and process, is it one of productivity, where the outcome from a business perspective is cost savings, or is it actually disruption? And in disruption, one has the power to rewire how one competes and change to a non-linear model. And the best example, again, is Blockbuster versus Netflix. They were in the same business. One went bankrupt and one became a FANG. And it was all down to asking the right question about disruption. But the key thing that's less well understood is they didn't just change the product and their mechanism, the entire organization, their set of processes, everything in culture changed. Let's just put it like this, Netflix didn't have a real estate team, as you'd imagine. And so that question then applies to the investment business within any of the firms that we serve. Are we swapping out, or does it fundamentally rewire? So there are some of the questions and observations we have from our side. So I'd like to pose the panel a question, which is, where do you think the prospects and connections you have underestimate what's currently possible with AI. And by with AI, I don't just mean the

foundational models. I mean AI with a purposeful use of a multi-agentic layer on top of an industrial scale.

## Bernie George

Yeah, that's a great, great experiment. So I think we think about those kind of the three curves you have. the foundation model capabilities, what you can accomplish with Agentic and other orchestrations on top of that, and then where you need to be to do human level work in an industry where there's essentially 0 tolerance for error. So with LLMs, you know, frontier models on the order of 50% hallucination rate. Tool calling is terrible. The accuracy, the efficiency, accuracy of tool calling is on the order of 50 to 70% on frontier models. So when you think about what's required to do human-level work, it's the chain of 10, 20, 30 steps, which need to be where error compounds. So if you have 50%, 75%, even 90%, or sorry, even 5% error on any one of those steps, and you compound that, you will not get a good result. So I think One of the key perceptions out there, I think, is that AI is moving fast, but it's not quite there to accomplish human-level work in our industry. But I think the reality is what we're seeing for many use cases is with the right orchestrations and frameworks and domain-specific processing of, you know, 20-page Excel models with 300 lines on each page or loan docs or, you know, transcripts or decks with all kinds of assorted, ways of presenting information with the right types of with the right types of orchestrations and analytical processes as well, you can create that human-level work. And if you can do that at the atomic level, then you can then chain that together into these pretty advanced analyst-associate BP level work product and do that, do that human-level work.

## Niall Boland

Yeah, thank you, party. I mean, for what it's worth, the way we think about it is you can, AGI was mentioned in one of the earlier panels, and what we've sort of seen is you can kind of hack near AGI if you have that industrial-scale agentic layer on top. And, you know, we've had analogies through all brands of technology from mobile telephony to hack around the limits of physics was to set up cells. right? That was the propagation hack, right? And we see that in everything. So I think that's something that's really poorly understood is that there's that gap to fill. So thank you, Bernie. Don.

## Don Muir

Yeah, 95% of enterprise pilots fail. I read the MIT study in my engineering team. We were kind of discussing it on site earlier this week, actually. 5% succeed. What separates the 5% from the 95%? We're a Y Combinator startup. We follow kind of the ideology of building a

partially, we call it partially autonomous AI experience for our customers. So we don't believe in end-to-end automation. We're not building a black box experience for our customers. We believe in guiding, providing tools for investment professionals to make them faster, better, stronger, so they can compete with their peers who are successfully adopting this technology. And so what does that come down to? Deep customization. So actually mirroring the workflows of investment professionals, meeting them where they're at. So replicating, creating templates to replicate their screening memos, their IC memos, their flash reports, their portfolio monitoring tools. So that firms across, you know, deal teams across level from analyst to MD, on the private credit side, they're consuming AI in a way that feels normal to them. That isn't too far off kind of the beaten path. Their analysts are now leveraging agents in talking to their agents as if they have their own dedicated analyst. And so what we're seeing with the customers we're deploying with is 2 things. One, we're customizing these reports in a way that mirrors their internal format. That includes tone, that includes context, so piping in deal repositories so we can actually contextualize underwrites and benchmark against deals they've done in the past. And then equipping each associate through VP with a dedicated analyst that they're actually interacting with to conduct asset level due diligence. So rather than giving them the answer and saying, here's the raw data, we're going to output a fully polished investment committee memo, we help guide them to that answer by leveraging them with the technology that allows them to move faster, which means they can kill bad deals faster out of box. That means they can focus their time and attention iterating with their team on the deals that are actually going to move the needle for their fund because it maps to their buy box and they've killed, kind of they've cut through all the noise much faster. And so At least with F2, what we've found success in is equipping deal teams with tools that make them more productive and more efficient while speaking the same language as them so they're not deviating too far from their existing workflows, their manual workflows that they've been using for the last three decades.

## Niall Boland

Thanks, John. Ready.

## Benji Cohen

I would say that the two big points from there that we've combined from both Don and Bernie are As Don was saying, everything needs to be hyper-specific. AI is not deterministic, as we've operated in a technology world before this. And so building something that is customized for, as somebody said on a couple of panels ago, a vertical slice is critically important. And then, as Bernie said, chaining these things together. You

don't actually have to solve everything all at once, but you can solve some things. You can then empower people to make smarter investment decisions. And so for me, I look back at our T-Rex experience where we were focused, hyper-focused on structured finance. We built this product that we called Quantum that was an LLM for 300-page offering memoranda that could effectively be modeled on T-Rex because we had this system that was just built for structured finance in 2 minutes. This used to take a week to two weeks of a very well-trained analyst's time to model in Excel. And you'd have 95% of it in 2 minutes. And you can then, that opens the range for the very smart people anywhere on Wall Street at Apollo, Blackstone, et cetera, to be able to go do this in private credit empowers them so much more. taking that to a very specific use, an even more specific use case with Camber Credit and putting it in Niall's context of being outcome focused, this gives us the leverage by chaining together all of these different processes, it gives us the leverage to go focus on much smaller portfolios that are not economic for for very large players. So there's roughly \$400 billion of the \$1.5 trillion US auto market is sitting on regional banks, credit unions, dealerships, balance sheets. It's not worth it for Blackstone with a \$250 billion dry powder for a private creditor thereabouts to go chase \$50 million portfolios or even \$250 million portfolios. They don't have the leverage that is this specific. We have built it there. And so we can go evaluate portfolios at scale, at speed, and get more yield out of them than they would otherwise and build an entire business out of that because we have taken an entire vertical slice. We have the benefit of being AI native. We don't have to go replace something else. And we can operate there again, efficiently and at scale and get more yield out of it.

**Niall Boland**

Cameron.

**Cameron Ladd**

Yep.

**Niall Boland**

Thanks, Ben.

**Cameron Ladd**

I'll echo what the previous panelists said. I mean, we've really used it as a supplement kind of to our existing workforce. Obviously, this is the current footprint that'll be probably totally different next year. But like right now, I think a use case we've started to mess around with is actually piping all this data through the back end of Snowflake. having like a group e-mail address that our ops users get, receive emails at, it'll run these basic queries through

Cortex and return basically a draft for our ops users to review because they're asking very pedestrian trade data and Snowflake would store that basic same data. So really we're thinking of it more of as a supplement to our existing users in this current state rather than a replacement and maybe just empowering our users and doing more with less.

## Niall Boland

Great, thank you very much. So if I were to sort of summarize what I'm hearing is that the question was, you know, where are people underestimating AI? It's that with a layer on top, you can mitigate some of the challenges that you get from just using an LLM. You can handle more scale, and therefore point it at use cases you previously couldn't point it at, and also kind of touch some of the tail, as you mentioned, Ben. And so you're starting to actually materially affect already today, just with that, some of the processes. So is that a fair reflection? Okay. I mentioned trying to describe the shape of the pool and kind of where the deep end is. So in some of our interactions, we interact with some of the sovereign wealth funds who are very well resourced, as you'd imagine, in this area. And so one of the more sophisticated deployments we've seen so far is where they're going far beyond just the sort of productivity use cases, right? And they actually, two of them now, have an AI member of the top-level executive committee voting on deals and asking questions of the deal team. And so that AI executive member has access not just to the memorandum, but to the entire deal room, and also to the broader set of information that the firm has accumulated and kind of rendered available. So in that particular case, and I kind of characterize that as the deep end, where it's not just making somebody more efficient, it's starting to change what the actual, the process is. They haven't changed the entire organization structure yet, but you can kind of see it on the com, I guess. And so one of the next questions I'd have for you is, you know, my definition of disruption is, you know, when the processes fundamentally change and the organization structure change and maybe even the products change. And so if the investment industry is based on a cognitive load premise that defines those currently, I'm curious if any of the panel or any of the audience for that matter have seen cases where already or there's a plan for a material change in what the business is or how it's structured on the basis of what's already possible.

## Bernie George

Yeah, I can speak to that a little bit. So I think there's sort of like 2 approaches. One is sort of the the high-level, disruptive, strategic push. And then two is sort of these use cases on the ground, these tactical, they save time, provide immediate ROI. And I think both of those seem like they're essential parts of the AI adoption strategy. At the deep end of the pool, at the disruptive end, I think There's, I think, at a high level, one way to think about it is

essentially the merging of quantitative and fundamental investing. historically, there have been quantitative investors who have taken structured data, portfolio alpha, alpha signals, portfolio construction, and then trade. And there have been fundamental investors who have taken unstructured data, have qualitative and quantitative thought processes around that, and then portfolio comes out. But now there's an opportunity to essentially merge those two information sets. And I think that's sort of the direction that a lot of these investment processes are evolving. But they're not going to, and they can't, because there are LP constraints on investment strategy, and you can't strategy drift, so they can't jump their right away. But if you can get confidence in that pyramid that leads up to the decision, so all the information that flows one layer at a time, that's sort of the end game. And I think we're a few layers deep. We're seeing it do the work of experienced analysts who paid 7 figures and And now parts of the entire parts of their job are just being done by AI. Done in parallel at first, but now just trusted. So I think, yeah, there's a very exciting end game. And yeah, different firms are at different stages of that evolution.

**Niall Boland**

Thanks, Barney. Ben, John, sorry.

**Don Muir**

Yeah, we see a pretty wide spectrum of adoption across our client base. And I'd say that the pace or the kind of percentage penetration of AI within private credit funds and banks tends to be inversely correlated to their size. So what we're seeing is in the lower middle market and middle market and direct lending, we're seeing widespread adoption. And we're actually seeing kind of budgets for new hires being delayed in favor, in many cases, of adopting an AI native operating system. And I think it speaks to the velocity of deals getting done. It speaks to the competition in the market and the fear that peers are developing a technological advantage in light of that competition. We're seeing workflows from screening through underwriting and portfolio monitoring within kind of lower middle market and middle market become pretty widespread. And as you move up market and start working with larger funds, and oftentimes, and certainly the large banks, like the super regional banks that we're now rolling out pilots with, adoption's a lot slower. Makes sense. There's an AI council. The MDs and the kind of senior folks are very hesitant to adopt the technology. The junior team members who are using it in investment banking or in school, they want the tech, but they can't. They're oftentimes kind of walled off from getting their hands on it because I think that some of the senior folks are more hesitant or more skeptical of adoption.

**Niall Boland**

So what would be the most sophisticated one where they've actually changed an organization or their process as a result of what they've achieved technologically?

## Don Muir

Yeah, and I don't know if scale is always correlated with sophistication in our experience, but it oftentimes is. In any case, we try to identify low-hanging fruit. And so we'll start really small. We'll start with automating a screening memo for some kind of cookie cutter direct lending strategy within the larger platform and identify the early adopters. So a set of users typically at the junior to mid-levels. So again, associate through VP on these deal teams who are really leaning into AI and proving out the use case, and then expanding usage across the firm, so in different strategies, and then ideally with the seniors, so getting kind of output in front of seniors. It's a longer curve. I think they're getting there, and they're seeing the power of this technology. It's kind of slowly but suddenly up market. But for the last year, we've really seen this large discrepancy in adoption between the earliest kind of highest velocity or smaller shops and the larger enterprise accounts that we're working with.

## Niall Boland

Great, thank you. Beth.

## Benji Cohen

I have a slightly different experience than that, but it is, but not to say it's wrong, it just depends on the use case. The parts of fixed income finance that I've worked in are highly regulated. And you saw that before from just a security perspective. Meeting the security requirements, some of them are archaic and not really relevant anymore, but you have to do it because these are regulated institutions and they sort of hit the lowest common denominator of all of the regulators that are there. So not directly relevant to AI, but where there is any sort of technological change, AI or otherwise, it has to comply with whatever the regulation is. And so you might have very eager adoption and flexibility with the more typically more junior staff. And the senior ones have the purview of understanding all of the different limitations and outcomes that they're trying to achieve. And so again, somebody on a panel, a couple panels ago said that you need actually the vertical slice of everybody in the organization, from senior to mid-level to junior and everybody in between, to buy into adoption at the same time, which I think can be driven from an outcome-based experience. And whether those are, as a vendor, as I used to be, with clients, or where you're developing it on your own. And I might take that back now to your original metaphor of Blockbuster and Netflix. And I think one of the most important things, whether you are in a shop like

Cameron's, where they're going and taking out the paper and replacing it by stitching together new processes, AI integrations, et cetera, or like Camber, where we can be AI, we have the benefit of being AI native because we started this year. It is a mentality. And it's a mentality that you can, from a leadership perspective, pervade across the entire firm, where it's not just the applications that are focused on yield, that are focused on the investment outcomes, the raison d'etre of the firm, but it's everything else. So what is the, I challenge our people daily, and they're not actually a lot of us now. It's very different than a year ago. But still, even the people who are seasoned portfolio managers, our CIO, how could we be doing this differently? How could you be going through your e-mail differently? Should you be using superhuman AI, for example, to go through your e-mail and get 10x efficiency and even just better responsiveness? It has nothing to do with our products. has nothing to do with directly with our investment thesis getting more yield out, evaluating more portfolios, et cetera. But ultimately it comes back to the mentality of being a technology oriented and an AI native investment manager. And I think that ultimately breeds the outcomes, you know, similarly to Netflix not having a real estate division. It wasn't actually core to the thesis, but It was part of the mentality.

## Niall Boland

Yeah, thank you. For what it's worth, our perspective is that an AI native investment manager will think about two key things differently. One is their investment process will not be a funnel. The funnel is a human hack on a cognitive load bandwidth limitation. Nearly all processes we see are funnels. as opposed to doing a deep dive in everything simultaneously. And then the other element of an AI native investment manager will be that you think very differently about sectors, about silos within your investment teams and your research teams, because silos are usually a reflection of the required for density. Nick.

## Bernie George

So if you can just repeat my question so it's picked up by other thing. So my question is for Benji. You know, when you are doing DQs, RFPs, whatever, and or investors are looking at your products. And obviously a selling point is that it is clean sheet, AI native, technology driven, all that stuff. What concerns are they? Are they expressing any concerns? Are there any special anxieties they have about that approach?

## Niall Boland

So the question, as I understand it, is whether clients have concerns about the approach which you take. on applying AI to DTQ processes given you're AI native? Is that? Yeah.

## Benji Cohen

Fantastic question. And I would say that it cuts both ways. So when technology feels like a black box when it is new, and we're talking to old school family offices, sovereign wealth funds, et cetera, they want to understand. And so there are a couple of ways to deal with it. You can present raw data and cleanse data. And so, you know, we're fully open kimono. here's the data, here's where it's come from, here's the process, we can spell it all the way out for you. are never going to go through all of this at scale. Also, we aren't. But go back and test anything. And so that's number one. However, I would put it in another context as well. So we were asked by an insurance-backed capital slash family office earlier this week For those who have been following anything in the news in the last couple of weeks in the auto industry, there was a massive news story about Tree Calore, which is one of the top 10 used car auto shops in the country, massive fraud, bankruptcy in deep subprime. And they asked us, how do you ensure that the data that you're getting, because at the end of the day, when you're buying used car loans, there are all sorts of assumptions, connotations, et cetera, of fraud for good reason. It happens not everywhere in the industry, not even in the majority of places, but in enough places. And you get big ones like this and it's highlighted. And so we walked through what our entire data process and how we have built AI to go look, to take unstructured data. Well, it's actually more, these are more our machine learning processes, but they're used somewhat synonymously and for an allocator like this, AI and machine learning are used effectively the same. And we walk them through, here's how we take data, here's how we take the unstructured data, make it structured so we can run our analytics, so we can build our investment portfolios out of it. And we go back and we iterate with whoever the seller of the portfolio is. And as we describe that process in great detail, they had extreme comfort that we almost couldn't, in a fail-safe way, fall prey to buying something like a tree color portfolio that would have fraud in it, because our AI is picking up the anomalies, even if it is not fraud. It's picking up, you know, we've got two different UUIDs for the same car. How is that possible? Maybe it's a different point in time. Maybe it was fat-fingered. There are a million different reasons. Maybe it is actually fraud, and you see a pattern of this, and they're ending up in two different same car. ending up in two different portfolios, which is a bit of what we had there. And so to answer Nick's question again directly, in summary, you can use AI defensively in a DDQ process.

## Cameron Ladd

I think that was my card.

## Niall Boland

So just one thing to add on that is, you know, I think it's an important extension is The promise of AI with the multi-agentic layer is not just to do things more efficiently, but it's actually to be able to see things that previously were not seeable, either because you couldn't get to that information or you couldn't join the dots and create the inference. And the analogy is a bit like 5th generation aircraft. It's called BVR, Beyond Visual Radar, where it's able to process things that the pilot couldn't process. The pilot's job is now a different job. It's actually defining the mission, it's not flying the plane, right? And it can fly at the edge of kind of aerospacial kind of stability. And so I think that's a key point to take away, is that we're kind of right at the point, and we've seen enough proof points, you've seen some, I've seen some, Bernie you've seen some as well, we all have, where the game starts to change when new things become possible beyond just efficiency. I'm conscious we're coming up to the end of the time. I just kind of want to open it up.

## Cameron Ladd

I just have a question for any of the panelists. So part of the investment process, you talked about AI and the portfolio construction and all that. On the back side, there's risk management. So what tools are in the toolkit to mitigate risk? Because once you've constructed your portfolio, that's great. And then you have to manage it and make sure that it's actually performing, or that any canary in a coal mine can be identified early on. So just a general question on that.

## Niall Boland

Bernie, you want to go?

## Bernie George

Yeah, I can speak to that a bit. So I think it's a great application. And I think part of it comes down to the general principle that AI can just process way more information than humans can. So things like, an 8K comes out, there's news, there's any, whatever, there's some subscriptions you have and there's some nugget in there. There's a court case filing and there's some nugget in there. Whatever it is, all of those can trigger multi-agentic processes, which can then evaluate and escalate as appropriate. And not just for public markets, but private markets as well. So quarterly reporting, consuming all the All the updates from portfolio companies, private credit, private equity, all of those, there's certainly canaries that are in the coal mine, which are not being detected with human processes and even with semi-structured processes now. So huge, huge opportunity there.

## Niall Boland

I'd add one thing as well. Thanks, Bernie, which is having been a user of risk tools as a hedge fund PM for a long time and and a good few gray hairs to come as a result as well. One great frustration I had for over a decade about any risk tool was, you know, they didn't really help me make decisions because usually they would just kind of outline the inventory of risks that I was exposed to, but it offered very little perspective on handicapping what the likelihood of those risks was. And there was no path to connect that information to the decisions I would have had to take, which is What's my pricing? What's my sizing? What's my hedging? What's my conditionality? And so the promise of AI on the risk side is to, again, see more things, as per Bernie's saying, but also then to have that kind of logical transmission layer to actually answer the so what question, which previously risk tools really, really didn't.

## Benji Cohen

Can I give one more example? And just also for Kudu AI Transcript Recorder, the question was, was how do you use risk management, or how do you pull AI into the risk management sphere after investment? For us, it is all the way downstream. I've never been in anything attached to consumer beforehand, but now we have auto loans and leases. And it's the most exciting piece for me post-investment, where the risk management and getting yield out of it is entirely about servicing. And so we're not quite there yet, but the next stage, which will start in the next three to six months, is building an AI native servicer. So instead of somebody, if you're in near prime or subprime, our focus is near prime, those mid-scale FICOs, vast majority of people pay their loans on time. But there is a portion that don't. And so then somebody is typically picking up the phone from a call center and calling them and running through a script. And I probably don't need to say more about how much better AI can do that. In Gen. Z and below, what we have seen is, or what we have seen from market research is that they would actually prefer to interact first with AI. They want AI texting them, AI that knows them, AI that's pulling a bunch of data about them already, knowing exactly how to speak to them, even when they're picking up on the phone. This is an AI agent speaking to you in an accent with a voice, using the words, using the cadence that speaks to the single mother of three in Kansas City driving a Ford F-150, 2023 Ford F-150, versus the 65-year-old retiree in Bakersfield, California, driving the exact same car. Completely different scripts for that. So you can imagine the psychology of what that does. But then you also get to interventions, which gets us back directly to risk management. So if somebody's missed a payment by 15 days and they've exhibited this, a bunch of other behavior before that, we know that is looking like it's going to go from a near-term delinquency to a long-term default. And there are interventions that we can feed through that AI directly or even have a human intervene that will create a greater recovery. And so there's actually not just much lower operational cost for us by not having all of those call

center people, but better outcomes on the investment side because of the AI intervention.  
Great.

**Niall Boland**

Thanks very much, Ben. I think we've sort of come to the end of the allotted time. So with that, I'll thank the panel. And again, thank you for inviting us all, Nick and team.